

Advancements in Particle Position and Size Determination for the Optical Measurement of Disperse Two Phase Flows under Optical Constraints

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. Christian Sax

Tag der mündlichen Prüfung:
Hauptreferentin:
Korreferent:
Korreferent:

26.09.2025
Prof. Dr.-Ing. Bettina Frohnapfel
Prof. Dr. Roland Griesmaier
Prof. DPhil Alfonso Castrejon-Pita



This document is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/deed.en>

Abstract

Disperse two-phase flows (DTPFs), comprising particles such as bubbles or droplets suspended in a continuous fluid phase, are central to numerous natural and industrial processes. Accurate, three-dimensional (3D) measurement of these flows is essential for advancing theoretical understanding and improving system efficiency and sustainability. However, conventional optical techniques often fall short in optically constrained environments, where multi-camera setups are infeasible and particle image (PI) overlap and low signal-to-noise ratios (SNR) pose significant challenges.

This thesis lays the groundwork for a novel framework enabling simultaneous 3D position and size determination of both tracer and dispersed phase particles using only a single optical access. Building on defocusing particle tracking velocimetry (DPTV) and interferometric particle imaging (IPI), the method extends IPI to the backscatter regime, making it suitable for constrained setups, and validates this extension experimentally. An inverse problem (IP) approach is developed to jointly estimate particle positions and sizes, enhancing robustness against noise and PI overlap. The method also enables the sizing of 3D-distributed bubbles or droplets and is experimentally validated against benchmarked methods. Convolutional neural networks (CNNs) are employed to distinguish between particle types, allowing two-phase measurements without additional hardware, and are tested across six different datasets. Additionally, a statistical model of PI overlap is introduced and validated through four experiments to guide experimental design and define operational limits. While each contribution addresses a distinct challenge in two-phase flow diagnostics, they collectively establish the foundational components for future integration into a comprehensive measurement framework.

The combined methodology aims to enable in-situ diagnostics of DTPFs in previously inaccessible environments, using standard PIV imaging equipment. It supports higher seeding densities, improved spatial resolution, and access to flow metrics such as slip velocities. The findings contribute valuable insights to experimental fluid mechanics, offering a cost-effective and scalable approach to complex flow diagnostics with potential benefits for both research and industrial applications.

Kurzfassung

Disperse Zweiphasenströmungen (DTPFs), bestehend aus Partikeln wie Blasen oder Tröpfchen, die in einer kontinuierlichen Fluidphase suspendiert sind, stehen im Zentrum zahlreicher natürlicher und industrieller Prozesse. Eine präzise, dreidimensionale (3D) Messung dieser Strömungen ist entscheidend für das Vorantreiben theoretischer Erkenntnisse sowie zur Verbesserung der Effizienz und Nachhaltigkeit von Systemen. Konventionelle optische Techniken stoßen jedoch häufig an ihre Grenzen in optisch eingeschränkten Umgebungen, in denen Mehrkamerasysteme nicht realisierbar sind und Überlagerungen von Partikelbildern (PI) sowie ein niedriges Signal-Rausch-Verhältnis (SNR) erhebliche Herausforderungen darstellen.

Diese Arbeit legt das Fundament für ein neuartiges Framework, das die gleichzeitige 3D-Positions- und Größenbestimmung sowohl von Tracer- als auch von dispergierten Phasenpartikeln mit nur einem einzigen optischen Zugang ermöglicht. Aufbauend auf der Defocusing Particle Tracking Velocimetry (DPTV) und der Interferometric Particle Imaging (IPI) Methode wird die IPI Methode erweitert, um im Rückstrebereich einsetzbar zu sein. Dadurch wird IPI für optisch eingeschränkte Systeme geeignet, und die Erweiterung wird experimentell validiert. Ein auf inversen Problemen (IP) basierender Ansatz wird entwickelt, um Partikelpositionen und -größen gemeinsam zu bestimmen, was die Robustheit gegenüber Rauschen und PI-Überlagerung erhöht. Die Methode ermöglicht zudem die Größenbestimmung von 3D-verteilten Blasen oder Tröpfchen und wird auf experimentellen Daten mit etablierten Methoden verglichen. Convolutional Neural Networks (CNNs) werden eingesetzt, um zwischen Partikeltypen zu unterscheiden, was Zweiphasenmessungen ohne zusätzliche Hardware erlaubt. Die Methode wird an sechs verschiedenen Datensätzen validiert. Zusätzlich wird ein statistisches Modell zur PI-Überlagerung eingeführt und in vier Experimenten validiert, um die Versuchsplanung zu unterstützen und Betriebsgrenzen zu definieren. Während jeder Beitrag eine spezifische Herausforderung in der Zweiphasenströmungsdiagnostik adressiert, bilden sie gemeinsam die grundlegenden Bausteine für eine zukünftige Integration in eine umfassende Messmethodik.

Die kombinierte Methodik zielt darauf ab, In-situ-Diagnostik von DTPFs in zuvor unzugänglichen Umgebungen mit Standard-Particle-Image-Velocimetry-Bildgebungsausrüstung zu ermöglichen. Sie unterstützt höhere Partikelkonzentration, verbesserte räumliche Auflösung und den Zugang zu Strömungskennwerten wie Schlupfgeschwindigkeiten. Die Ergebnisse leisten einen wertvollen Beitrag zur experimentellen Fluidmechanik und eröffnen einen kosteneffizienten und skalierbaren Ansatz für die komplexe Strömungsdiagnostik mit vielversprechendem Potenzial für Forschungs- und Industrieanwendungen.

Acknowledgements

First and foremost, I would like to thank Bettina Frohnapfel for the opportunity to pursue this thesis under her supervision and within her research group. I also extend my thanks to Jochen Kriegseis, who acted as my direct supervisor, for sharing his expertise and his guidance throughout this time.

I would like to thank MathSEE, the Karlsruhe Graduate School of Computer and Data Science (KCDS), and especially Angela Hühnerfuß, whose support made this journey possible. The interdisciplinary nature of this PhD thesis would not have been achievable without the unique framework and opportunities provided by KCDS.

A very special thanks goes to Roland Griesmaier, whose steadfast support and mentorship were instrumental throughout this journey. His guidance extended far beyond the academic scope of this thesis, and I am deeply grateful for the trust, encouragement, and opportunities he provided during a pivotal time in my career.

I am deeply grateful to Alfonso Castrejon-Pita for welcoming me into his research group at the University of Oxford and for the opportunity to spend seven months immersed in such a vibrant and inspiring academic environment. This experience was not only a highlight of my doctoral journey but also a deeply formative chapter in my life. His generosity and support left a lasting impact, both professionally and personally.

I want to thank to Alexander Wagner and Alexander Kaufmann for their generous advice and unwavering support throughout my doctoral journey. Although not formally involved in this thesis, their guidance and encouragement came at crucial moments, and their belief in my potential made a lasting difference.

I am deeply grateful to my family for their unwavering support and for standing by me through every challenge along the way. Their encouragement has been a constant source of strength.

I am also thankful for my friends and the incredible people I've met over the past three years - their friendship, kindness, and shared moments have made this journey all the more meaningful. I further want to thank my cat, Sushi, for always bringing joy to my life and for volunteering as the photo model for Fig. 2.19.

I gratefully acknowledge the financial support from the Karlsruhe House of Young Scientists (KHYS). The research visit at the Fluid Mechanics Lab at the University of Oxford was partially funded by the Karlsruhe House of Young Scientists (KHYS).

This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) via Project Grant KR4775/4-1 within the Research Unit FOR 5595 Archimedes - Project Number 510921053.

I want to furthermore thank Leister *et al.*, [115], Pasch *et al.*, [149] and Pöppe [151] for kindly providing experimental raw images as used in this work.

Contents

Abstract	i
Kurzfassung	iii
Acknowledgements	v
1 Introduction	1
1.1 Motivation	2
1.2 Research objective and outline of the study	4
2 Fundamentals	7
2.1 Two-Phase Flows	7
2.1.1 Definitions and Classifications	7
2.1.2 Flow Patterns and Flow Regimes of Two-Phase Flows	8
2.2 Selected Measurement Techniques for the Investigation of Disperse Two-Phase Flows	10
2.2.1 Overview: Measurement Techniques for Flow Velocity and Particle Size . .	10
2.2.2 Defocusing Particle Tracking Velocimetry	12
2.2.3 Interferometric Particle Imaging	18
2.3 Relevant Concepts in Optics	21
2.3.1 Light Scattering at Droplets and Bubbles	21
2.3.2 Ray Transfer Matrix Model	28
2.3.3 Propagation of Light waves	30
2.4 Relevant Basics in Deep Learning for Particle Detection, Classification and	
Image Generation	33
2.4.1 Neural Networks	33
2.4.2 Training of Neural Networks	35
2.4.3 Convolutional Neural Networks	38
2.4.4 Object Detection Networks	40
2.4.5 Generative Adversarial Networks	45
3 Interferometric Particle Imaging for Particle Sizing in Front-, Side- and Backscatter	47
3.1 An alternative Transfer Function for Interferometric Particle Imaging	47
3.2 Theoretical Model for IPI in Front- Side- and Backscatter	49
3.2.1 Visibility of Interference Patterns	49
3.2.2 Visibility Plots and Examples	51
3.2.3 Determination of the Glare Point Spacing	53
3.3 Experimental Validation of Backscatter IPI	57
3.3.1 Experimental Procedure	57
3.3.2 Data Processing	59
3.4 Uncertainties and Limitations of IPI	60
3.4.1 Sources of Uncertainty	60
3.4.2 Main Influences on the Uncertainty	62
3.5 Concluding Remarks	66
4 Inverse Problem Approach for the Determination of Particle Location and Size . .	69

4.1	Combination of Defocusing Particle Tracking and Interferometric Particle Imaging	69
4.2	The Forward Model - Direct Problem	70
4.2.1	Scattering Model	70
4.2.2	Wave Propagation Model	72
4.2.3	Validation of the Forward Model	75
4.3	Inverse Problem for a Single Particle	78
4.3.1	Mismatch Functionals for the Optimization	78
4.3.2	Optimization Strategy for a Single Particle	80
4.3.3	Test on a Single Synthetic Particle Image	81
4.3.4	Transferability to Real Particle Images and Robustness of IP Approach . .	85
4.4	Inverse Problem for Multiple Particles	87
4.4.1	Mismatch Functionals and Optimization	87
4.4.2	Tests on Synthetic Data	88
4.4.3	Particle Image Overlap	90
4.4.4	Particle Reconstruction from the IP Approach	91
4.5	Discussion and Conclusions	92
5	Simultaneous Measurement of two Phases in Defocusing Particle Tracking Velocimetry	95
5.1	Introduction	95
5.2	Theoretical Concept of Pattern based Phase Distinction in Defocused Images . . .	97
5.3	Auto-labeling Approach and Training of Networks for Particle Detection and Classification	98
5.3.1	Conceptual Approach for the Generation of Auto-labeled Training Data . .	100
5.3.2	Training the Generative Adversarial Networks for Image Generation	101
5.3.3	Training of Object Detection Networks	103
5.4	Detection and Separation of Different Particle Image Types	104
5.4.1	Test Dataset	104
5.4.2	Evaluation Metrics	106
5.4.3	Detection Results	108
5.5	Concluding Remarks	110
6	On the Particle Image Overlap in Single Camera Defocusing Approaches without Astigmatism	113
6.1	Introduction	113
6.2	Datasets for the Empirical Investigation	114
6.2.1	Generation of a Dataset for an Empirical Investigation	114
6.2.2	Experimental Datasets for Validation	116
6.3	Definition of Overlap Metrics	117
6.3.1	Metrics for the Number of Overlaps	117
6.3.2	Metrics for the Amount of Overlaps	118
6.3.3	Direct Computation of Overlapped Particle Image Fraction	121
6.3.4	Definition of a Scaling Metric	122
6.4	Scaling Model for Particle Image Overlap	123
6.5	Concluding Remarks on Particle Image Overlap	126
7	Conclusions and Outlook	129
8	References	135
	Bibliography	135
	List of Conference Contributions	153
	List of Journal Publications	154
	Acronyms and symbols	155

List of Figures	165
List of Tables	175
A Appendix: Interferometric Particle Imaging in Front-, Side- and Backscatter	177
A.1 Further Information on the Optical System of the Validation Experiment	177
A.2 Algorithm for the Calculation of Glare Point Maps	178
B Appendix: Neural Networks and Training for Two-Phase DPTV	181
B.1 GAN Training Configuration and Training Scores	181
B.2 Training Losses and Configuration for the Object Detection Networks	182
B.3 False Positive Class Bias	183
C Appendix: Particle Image Overlap in Single Camera Defocusing Approaches without Astigmatism	185

1 Introduction

Multiphase flows are ubiquitous, ranging from vast scales such as interstellar clouds in the far reaches of the Universe[192], through weather phenomena like rain and snowfall, to the sparkling texture of beverages in our daily lives. A multiphase flow is any fluid flow comprising more than one thermodynamic phase, separated at a scale well above the molecular level[21]. A special category of multiphase flows is disperse two-phase flows (DTPF), which consist of finite particles (the dispersed phase, e.g. bubbles or droplets) suspended in a continuous phase (the carrier phase, e.g. water or air) forming a connected volume[21].

DTPFs are of key interest in many research fields, including geology, physics, chemical and mechanical engineering, and even the food industry. Geophysical phenomena such as volcanic eruptions produce winds carrying particles like rocks or ash, which can lead to subsequent weather phenomena. Weather events such as fog, rain, snowfall, and dust storms are of particular interest to environmental research [11], especially in the context of climate change, where a better understanding of these processes is crucial. Such airflows, which contain particles such as ice and water droplets, have critical implications for air travel. For example, they affect the icing of aerofoils, which not only reduces the fuel efficiency of aircraft but can also pose a major safety concern for thousands of people flying every day. However, aircraft must contend with DTPFs not only on their wings but also within their engines. Most fossil-fuel-based modes of transport, from aircraft to cars and ships, encounter DTPFs in combustion processes, such as fuel sprays in turbines and piston engines [25]. In these areas of research, DTPFs influence combustion efficiency through heat transfer and mixing behaviour, which in turn has a direct impact on fuel efficiency and engine longevity. Fuel efficiency and longevity offer clear advantages for sustainability and provide financial benefits for the industry. Progress in this area can therefore give vehicle and engine manufacturers a competitive edge. DTPFs are also present in power plants, such as gas and steam turbines, as well as in the oil and gas industry [219, 171]. For example, in offshore production, oil is extracted from the seabed as a mixture of crude oil, seawater, sand, and gas, which can pose hazards to expensive equipment [219]. Bringing oil up from the seabed or transporting other fluids requires pumps, which can suffer damage and inefficiencies due to cavitation, once again, a phenomenon involving DTPFs. When transporting fluids through pipelines, cavitation can lead to pressure loss, increasing the energy required to move the fluid and potentially even damaging pumps and pipelines. This can be a significant source of cost and energy waste in the oil and chemical industries. The chemical industry is concerned with DTPFs not only because of fluid transport in pipes, but perhaps more importantly due to their role in chemical reactions [219, 11]. Different phases in DTPFs can be products of chemical reactions. DTPFs also influence heat transfer and mixing behaviour, which in turn affect reaction dynamics. Understanding these phenomena can lead to more efficient processes, impact scalability, or even enable new processes to realise specific reactions. This is, without doubt, a highly desirable outcome in these fields. Another influence of DTPFs in chemical processes can be found in fuel cells, where bubble flow affects performance by reducing conductivity [183, 75]. Improving the efficiency of fuel cells can contribute to the advancement of climate-neutral buses, cars, and trains, as well as sustainable power generation. Another market that is heavily influenced by environmental changes is agriculture, where DTPFs also appear in the form of sprays used to protect crops from disease and certain insects. Improved

spraying methods help to reduce labour and the amount of wasted chemicals, thereby lowering both the cost and environmental impact of crop protection [55]. DTPFs occur not only in food production but also in food consumption. Bubbles in beverages affect both the texture and quality of the drink [65, 45]. The stability of foam and the perceived texture caused by rising bubbles play a crucial role in the quality of lemonades, soft drinks, beer, and sparkling wine. Strictly speaking, even tap water is a three-phase flow consisting of water, air, and microscopic particles. When such tap water is used to cook pasta, the boiling process creates a two-phase flow once again.

As phenomena involving DTPFs are ubiquitous, their study is of great interest—whether to help save our planet, gain a competitive advantage in business, or simply enjoy a better drink with friends. Understanding the underlying phenomena often requires experimental investigation. Numerical models must be calibrated and validated against experimental data. Physical models and theories need to be tested using real-world observations. Moreover, new phenomena can be discovered and verified through experiments, thereby extending existing models and deepening the general understanding of flow physics. It is therefore crucial to have experimental techniques capable of capturing such data. Among the various available methods, optical measurement techniques offer the advantage of being non-intrusive.

In optical measurements of DTPFs, the flow field of each phase and the size of the particles of the dispersed phase are often of primary interest. The particle size of the dispersed phase is a primary quantity that can be measured directly. In contrast, flow field-related information can only be determined indirectly from the positions of tracer particles in the flow over time. Since both phases are of interest, the positions of both tracer particles in the continuous phase and particles in the dispersed phase must be measured. Other primary quantities, such as pressure and temperature, may also be observed but are beyond the scope of this work. Therefore, the positions of tracer particles in the continuous phase, along with the positions and sizes of particles in the dispersed phase, are of particular interest, as they form the basis for a variety of secondary or derived quantities. Measuring the position of particles over time allows the derivation of flow velocities and the flow topology of each phase. When the velocities of both phases are measured simultaneously, slip velocities can be determined, which influence, for example, the lift forces acting on particles [85] and heat transfer [41]. Particle positions can also be used to describe mixing and sedimentation behaviour, which is crucial in many industrial, chemical, and environmental processes. While the size of dispersed phase particles is often of direct interest, such as the size distribution of droplets produced by a nozzle, other quantities can also be derived from particle sizes. Knowing the size of nearly all particles in the fluid allows the void fraction to be estimated. By measuring particle sizes over time, quantities such as bubble growth or droplet dispersion can be determined. As many quantities are derived from the primary measurements of position and size, it is important that these are determined with high accuracy initially, in order to minimise error propagation into derived metrics.

1.1 Motivation

While there are other methods for measuring DTPFs [150], optical measurement techniques are often more convenient and require less equipment. These techniques can be applied to DTPFs with a transparent continuous phase and offer the advantage of being non-intrusive, thereby not affecting the observed phenomena. However, the complexity of such measurements increases significantly with the dimensionality of the measurement, the spatial resolution relative to the measurement volume, restrictions on optical access, and the requirement to measure both phases simultaneously. Optical methods such as laser Doppler velocimetry (LDV) [5] allow for one-dimensional (1D) point

measurements. Techniques like particle imaging velocimetry (PIV) enable the measurement of flow velocities in two-dimensional (2D) planes illuminated by a light sheet [156]. The measurement becomes more complex when extended to 3D volumes. Techniques such as tomographic PTV [143, 173, 174] address this challenge but require multiple viewing angles to reconstruct the 3D positions of particles. PIV and PTV retrieve the fluid motion based on the displacement of tracer particles

Even greater complexity arises when boundary conditions limit optical access, effectively ruling out multi-camera techniques. Such constraints typically occur when measurements are taken in confined spaces, such as industrial machinery, e.g. inside a compressor or turbine chamber, or in environmental settings like small cavities or nooks. Some engineering applications do not permit a second optical access, as components essential to the machine's function would need to be removed to make space for it. Techniques such as holographic PTV [24] can reconstruct a three-component (3C) flow in 3D using only a single camera. However, holographic methods still require a second optical access to introduce light into the measurement volume. For three-dimensional, three-component (3D-3C) flow reconstruction with only a single optical access, the two commonly used methods are astigmatism PTV (APTV) [87] and defocusing PTV (DPTV) [215, 57]. These methods, however, face additional challenges due to intensity loss during defocusing, which results in low signal-to-noise ratios (SNR), and particle image (PI) overlap, which causes information to be incomplete or superimposed with other data and lower accuracy in the out-of-plane direction.

All of these methods are capable of measuring the movement of a single type of particle, i.e. single-phase flows. If a second phase is to be measured simultaneously, a combination of methods is often required, adding further complexity. To measure the flow field of both phases, both the tracer particles in the carrier phase and the particles of the dispersed phase must not only be tracked but also distinguished from one another. Particle separation is often achieved via the wavelength of the light emitted by the particles [97, 29], but this again requires additional equipment such as RGB cameras, wavelength filters, and multiple cameras. A measurement method for two-phase cases should combine the ability of DPTV and APTV to determine particle positions using only a single optical access with the capability to distinguish between different particle types. Additionally, such methods should be able to cope with low SNR resulting from defocusing and with the effects of PI overlap.

Another challenge is the sizing of dispersed phase particles. Techniques such as interferometric particle imaging (IPI) [98, 64, 144], digital in-line holography [102], and imaging can be used to size particles [53]. Among these techniques, IPI offers better size resolution relative to the measurement volume compared to both holography and imaging. Holography, on the other hand, can measure particles in 3D volumes, whereas IPI and imaging techniques are mostly restricted to 2D planes. However, both IPI and in-line holography typically require a second optical access to introduce light into the measurement volume, making them difficult to use in single optical access scenarios. Imaging can be used with a single optical access, but it is limited to quasi-2D measurement areas due to the restricted focal length of the imaging device, and it provides poorer size resolution compared to IPI. Ideally, a measurement technique would combine the size resolution of IPI with the single optical access capability of imaging and the ability of holography to size particles distributed in 3D volumes.

The simultaneous measurement of both phases with 3D-3C information and sufficient spatial resolution, along with the concurrent sizing of dispersed phase particles within this 3D volume, already presents a considerable challenge. When the additional constraint of single optical access is introduced, most existing measurement strategies fall short of meeting all these requirements. Nevertheless, DTPFs are inherently three-dimensional and non-stationary, which necessitates the

simultaneous measurement of both phases in 3D. This, in turn, imposes a range of stringent demands on measurement strategies in order to obtain the comprehensive information required for meaningful progress in the respective fields.

1.2 Research objective and outline of the study

This thesis investigates ways to extend DPTV and IPI by enhancing the underlying physical models and applying optimisation and deep learning methods for data processing. The goal is to derive information on the 3D positions and sizes of different types of particles in scenarios with limited optical access. Two-phase flows can exhibit a variety of flow regimes, not all of which can be captured using a single measurement technique. The scope of the present work is therefore limited to disperse two-phase flows with quasi-spherical particles, such as bubble flow and mist flow. The necessity of a single optical access measurement approach, capable of determining the 3D positions of both tracer particles and particles of the dispersed phase, while also sizing the dispersed phase particles, was established in the previous section. The development of such an approach can be broken down into smaller steps. First, a sizing method for the dispersed phase, namely IPI, which is also suitable for the combination with 3D position determination, must be adapted for evaluation in single optical access scenarios, i.e. in backscatter. The necessary changes in measurement evaluation due to the backscatter regime must be identified. In the next step, a method must be developed that can simultaneously determine particle size and 3D position. This approach must overcome specific challenges, such as the dependence of diameter estimates on particle position estimation, and issues arising from PI overlap. Initially, this is considered only for the particles of the dispersed phase, since tracer particles do not require sizing. The subsequent step extends the approach to include the tracer particles, enabling measurement of both phases within the DTPF. Once a method fulfilling the aforementioned requirements is established, practical guidance must be provided to the experimenter on how to set up an experiment in accordance with the limitations of such an approach. Therefore, the following research questions are formulated:

Research questions:

- RQ1 How can IPI be expanded from the front- to the side-scatter and backscatter regime to enable single optical access IPI?
- RQ2 How do uncertainties and limitations change when performing IPI in side- and backscatter compared to the front-scatter regime?
- RQ3 How can an inverse problem approach improve the determination of particle position and size?
- RQ4 Can neural networks be trained to reliably detect and differentiate tracer and dispersed phase particles in DPTV images, enabling simultaneous two-phase measurement?
- RQ5 How can the expected PI overlap in DPTV and IPI be quantified, as a function of experimental settings?

Research Questions (RQ1) and (RQ2) focus on extending the physical model used for IPI in order to exploit its high resolution compared to other techniques, while enabling single optical access capability. IPI is primarily employed in front-scatter configurations, where the light source and camera face each other with an angle, and in some cases, it is used in side-scatter arrangements.

However, current evaluation strategies are not applicable in backscatter, which is essential for single optical access scenarios. RQ1 addresses this limitation by developing a method to derive transfer functions between the interference pattern and particle size, independent of the scattering regime. Given that scattering properties differ in backscatter compared to front-scatter, RQ2 investigates the resulting effects on SNR, measurement accuracy, and other limitations that arise when performing IPI in backscatter, as opposed to the well-studied front-scatter regime. Both RQ1 and RQ2 are explored in Chapter 3. This chapter details how to evaluate IPI measurements in backscatter, identifies additional factors that must be considered, and explains how measurement uncertainties and limitations vary across front-, side-, and backscatter configurations. The content of Chapter 3 is based on Sax *et al.* [SDK25].

Research Question three (RQ3) then shifts focus to data evaluation, aiming to extract information about particle position and size. An optimisation method based on an inverse problem approach is developed to determine particle position and size from IPI images. Particular attention is given to overcoming the limitation that IPI is predominantly restricted to 2D planes, rather than enabling full 3D measurement volumes. Additionally, the issue of PI overlap, which compromises the accuracy of position and size determination in both IPI and DPTV, is addressed. Chapter 4 outlines the construction of a physical model and a mismatch functional for this inverse problem approach. It then explores how particle position and size estimation can be enhanced by employing the optimisation method, in contrast to state-of-the-art techniques currently used in IPI and DPTV. Additional challenges, such as low signal-to-noise ratio (SNR) and information loss due to PI overlap, are also addressed. The content of Chapter 4 is based on Sax *et al.* [SGK25].

With a method established to acquire data in backscatter and to evaluate particle positions and sizes simultaneously for a single particle type, the next research question (RQ4) addresses how to measure two phases simultaneously. Chapter 5 focuses on the detection and classification of two types of particles in DPTV, namely tracers in the continuous phase and dispersed phase particles. This classification can then be used either to track the phases directly using traditional approaches, or to initialise an inverse problem approach as described in Chapter 4. The content of Chapter 5 is based on Sax and Kriegseis [SK25b].

The final research question (RQ5) bridges the gap between the theoretical limitations of data evaluation and the practical acquisition of raw data. A statistical investigation of PI overlap offers insights into its nature and the development of a statistical model on PI overlap provides direct guidance to experimentalists on how to configure their experiments. This model establishes a clear link between algorithmic limitations related to PI overlap and the corresponding experimental settings. RQ5 is addressed in Chapter 6. The content of Chapter 6 is based on Sax and Kriegseis [SK25a].

While each research question targets a distinct technical challenge, they are unified by the overarching goal of enabling single-access, 3D, two-phase flow diagnostics. Rather than presenting a fully integrated measurement system, this thesis establishes the essential building blocks, validated independently, that enable such a system to be realized in future work.

The novelty of this thesis lies in setting the foundation for a single optical access measurement technique that combines IPI and DPTV for simultaneous 3D position and size determination of tracers and bubbles or droplets in DTPFs. Unlike existing methods, this approach enables backscatter-based measurements, integrates inverse problem optimisation for improved accuracy, and leverages deep learning for phase differentiation, all within a constrained optical setup. Furthermore, a novel statistical framework is introduced to quantify PI overlap as a function of experimental parameters, offering practical guidance for experiment design and ensuring data quality that aligns with the

limitations of post-processing in high-density flow conditions. By enabling accurate 3D measurements of both tracer and dispersed phase particles in dense, optically constrained environments, this work opens new possibilities for in-situ diagnostics of DTPFs in harder-to-access applications across a variety of research fields. This expanded accessibility not only advances the scientific understanding of multiphase flow dynamics but also supports the development of more efficient and sustainable technologies.

Two distinct citation styles are employed in this thesis. An alphabetical style (e.g. [ABC12]) is used to highlight publications authored by the candidate during the course of the PhD, distinguishing them from external sources. A numerical style (e.g. [123]) is used for referencing all other literature. Portions of this thesis have been published by the author or are currently under submission. These sections are marked with quotation marks and a disclaimer at the beginning of the relevant chapter or section to clearly indicate self-citations.

2 Fundamentals

This chapter aims to support the reader in understanding the following chapters by providing the necessary fundamentals and concepts referenced in Chapters 3, 4, 5, and 6. First, an introduction to multiphase flows is presented. In particular, dispersed two-phase flows, the focus of the measurement strategies developed in this work, are contextualised within the broader category of multiphase flows. A brief overview of measurement techniques for dispersed multiphase flows is then provided. Second, the two primary measurement techniques, DPTV and IPI, are explained in greater detail. These techniques form the foundation for the development of the measurement strategies discussed in this thesis. Third, basic optical concepts relevant to the development of these measurement strategies are introduced. This includes explanations of light propagation and light scattering by particles of relevant size. Finally, an overview of specific machine learning methods is given, as these are integral to the data evaluation methods developed in Chapter 5.

2.1 Two-Phase Flows

This section begins with an overview of multiphase flows, including their definitions and classifications. The second part focuses specifically on gas-liquid two-phase flows, such as air-water systems. Finally, DTPFs, the primary subject of measurement in this thesis, are contextualised within the broader spectrum of two-phase flow types.

2.1.1 Definitions and Classifications

A multiphase flow is defined as a flow consisting of two or more thermodynamic phases that are distinct well above the molecular level [21]. The simplest example is a two-phase flow, which constitutes the type of flow relevant to this treatise. However, more complex flows may involve three or more phases. There are many different ways to categorise multiphase flows. This section provides a brief overview of various types and classification methods. Perhaps the most straightforward method of categorising multiphase flows is by the state of aggregation. Two-phase flows can be divided into three categories: gas-liquid (e.g. bubbly flow), liquid-solid (e.g. sand particles in water), and gas-solid (e.g. hail in air). Three-phase flows may then include combinations such as gas-liquid-solid, and so on. It should be noted that a *phase* refers to a thermodynamic state such as gas, liquid, or solid. Chemically distinct substances, referred to as *components*, are excluded from this definition. For example, an oil-water mixture consists of two components but constitutes only a single phase. Another important distinction is that of *flow fields*, which are topologically distinct regions within a flow. For instance, an annular flow (see Section 2.1.2) features a water film along the boundaries and water droplets in the centre. Thus, the same two-phase flow contains two different flow fields [219]. It quickly becomes apparent that classification by state of aggregation alone is not sufficiently exhausting. For example, bubbles in water behave fundamentally differently from water droplets in air, despite both being liquid-gas flows. Additionally, different flow fields can significantly alter the behaviour of the flow. Another classification is the distinction between *dispersed* and *separated* flows [21]. Dispersed flows consist of finite particles (i.e. the

dispersed phase) suspended within a continuous carrier phase (i.e. the continuous phase). A typical characteristic of dispersed multiphase flows is that the carrier phase comprises the majority of the volume. Separated flows, on the other hand, contain two or more continuous streams divided by an interface, the boundary between two phases. It should be noted that partially separated flows also exist. For example, annular flow is both partially separated and partially dispersed [21]. As such, there exists a degree of separation (global vs local), making these two categories a spectrum rather than strictly distinct classes. Beyond the basic distinction between dispersed and separated flows, further sub-classifications are often necessary to describe the internal structure of the dispersed phase. For instance, within dispersed two-phase flows, one can distinguish between *homogeneous* and *inhomogeneous* particle distributions [219]. This distinction refers to whether the dispersed phase is uniformly distributed throughout the continuous phase or exhibits spatial variations (e.g. vertical bubbly flow due to the effect of gravity), which can significantly influence flow behaviour and measurement strategies. Another distinction is based on the presence of heat transfer, classifying flows as either *adiabatic* or *diabatic* [219]. This differentiation is important, as heat transfer has significant implications for phase changes such as vaporisation or condensation. Further classifications consider the direction of flow of each phase. Two-phase flows can be either *co-current* or *counter-current*, depending on whether the two phases flow in the same or opposite directions, respectively [219]. While the previous classifications focus on the physical arrangement and interaction of phases, another widely used approach categorises two-phase flows based on their observable geometric configurations of the interfaces, also referred to as *flow regimes* [21, 219]. The simplest examples involve quasi-one-dimensional pipe flows. For such flows, flow regime maps can be constructed. These maps typically relate the mass or volume flow rate of each phase and are generally valid only for a specific fluid pairing and pipe diameter [219]. In these cases, co-current flows in either vertical or horizontal pipes are typically investigated, which are discussed in more detail in the next section.

In summary, many different approaches exist to classify two-phase or even multiphase flows, and the most appropriate method often depends on the objective of the investigation.

2.1.2 Flow Patterns and Flow Regimes of Two-Phase Flows

As mentioned in the previous section, pipe flows are commonly used to illustrate different types of flow regimes. In such cases, flows are typically categorised as horizontal, vertical, or inclined. This distinction is particularly important when gravitational effects are significant, which is usually the case when there is a large density difference between the two phases (e.g. in air–water flows, the density ratio is approximately 1000) [219]. It should be noted that flow regime maps are generally not universal, and extrapolating results beyond a specific pipe diameter or fluid pairing is challenging [21]. Accordingly, this thesis employs flow regimes solely to contextualise DTPFs within the broader spectrum of flow types.

Flow Regimes in Vertical Pipe Flows

Flow regimes of co-current flows in vertical and inclined pipes typically include bubbly flow, slug flow, churn flow, annular flow, and dispersed (or mist) flow [209]. An overview of these flow regimes is given in Fig. 2.1. It should be noted that the boundaries between these regimes are not sharply defined but rather represent gradual transitions, which can make it difficult to clearly identify the prevailing flow regime in some cases. At low void fractions and low flow rates, bubble flow occurs. In this regime, the liquid serves as the carrier phase and the gas as the dispersed phase. As the flow rate increases, the gas bubbles become more disturbed and may break up into smaller, more

uniformly distributed bubbles. This leads to what is often termed bubbly flow, characterised by a homogeneous distribution of small gas bubbles within the liquid [219]. Some literature distinguishes between bubble and bubbly flow based on the degree of dispersion, with bubbly flow representing a more advanced stage of dispersion [219], whereas other sources refer to both simply as bubbly flow. Slug flow appears at higher void fractions, where gas bubbles coalesce to form larger bubbles. When these bubbles reach sizes comparable to the pipe diameter, they transition into plugs [219]. With further increases in flow velocity, slugs and plugs break down into an unstable regime known as churn flow. This regime represents the transition between slug or plug flow and annular flow [219]. In annular flow, the liquid forms a film along the pipe walls, while the gas flows through a central core. The liquid film may contain bubbles, and the gas stream in the centre can carry droplets [219]. Annular flow typically occurs at high gas volume fractions, with the gas phase moving faster than the liquid film. A variation known as wispy-annular flow includes larger liquid portions entrained in the central gas core. At very high gas volume fractions, the flow transitions into a fully dispersed droplet regime, also referred to as mist flow. In this case, the gas becomes the carrier phase, and the liquid exists as fine droplets. The gas core in annular flow can also be considered a form of mist flow.




Figure 2.1: Flow patterns in a vertical pipe. The flow patterns are also marked by whether they classify as a disperse or partially disperse partially continuous two-phase flow. Figure based on [21].

Flow Regimes in Horizontal Pipe Flows

Two-phase flow in horizontal pipes differs from that in vertical pipes, as gravity tends to stratify the flow [219]. Frequently used flow regime maps and classifications for horizontal pipe flows were developed by Baker [10] and Weisman [210]. An overview of the flow regimes in horizontal pipe flows is shown in Fig. 2.2. Bubbly flows in horizontal pipes are similar to those in vertical pipes, characterised by dispersed bubbles and low void fractions. However, due to the lift force acting on the bubbles, they are usually heterogeneously distributed across the pipe cross-section. Stratified flows occur when the two phases are fully separated, with liquid at the bottom and gas at the top. These flows can be either stratified smooth or stratified wavy, depending on the interface behaviour [219]. Annular flows also occur in horizontal pipes. The main difference compared to vertical pipes is that the liquid film is typically thicker at the bottom and thinner at the top due to gravitational effects. Intermittent flows, such as plug and slug flow, also appear in horizontal

configurations. In these regimes, gas bubbles are elongated and differ from the nearly spherical bubbles observed in bubbly flow.




Figure 2.2: Flow patterns in a horizontal pipe. The flow patterns are marked by whether they classify as a disperse, partially disperse partially continuous or continuous two-phase flow. Figure based on [21].

2.2 Selected Measurement Techniques for the Investigation of Disperse Two-Phase Flows

As demonstrated in the previous section, there are many different kinds of multiphase flows. Consequently, a variety of measurement methods have been developed to characterise these different types of flows. Different techniques pursue different objectives: some aim to identify the flow regime, while others are specialised in deriving specific quantities for a given regime. This section introduces the two main measurement techniques, employed throughout this thesis, for the determination of the primary quantities particle position and size. The techniques and concepts presented here are applied in Chapters 3, 4, 5, and 6.

2.2.1 Overview: Measurement Techniques for Flow Velocity and Particle Size

In dispersed two-phase flows, several primary quantities are of interest. These include the velocities of the phases, namely, the velocity field of the carrier phase and the motion of individual dispersed

particles, as well as the sizes of the dispersed particles, and other quantities such as pressure and temperature. As this thesis focuses on the velocities of both phases and the sizes of the dispersed particles, only measurement methods relevant to these quantities are considered. Techniques for measuring other parameters, such as pressure and temperature, do exist but fall outside the scope of this treatise. Furthermore, only non-invasive methods are considered within the scope of this work.

Determining Velocity

There are multiple techniques available to measure flow velocity, which can be applied to either the continuous or dispersed phase, and in some cases, both. Electrical impedance tomography (EIT) [54, 26] can be used to measure volumetrically averaged velocities. EIT detects changes in impedance caused by the differing electrical properties of the phases in a two-phase flow. It can also be used to determine quantities such as void fraction or to identify the flow regime [28]. However, EIT cannot resolve complex three-dimensional velocity fields [150]. X-ray radiography can be employed to measure particle motion, particularly in high-density suspensions [138] or when the continuous phase is opaque, rendering optical methods ineffective. A related technique, neutron radiology, is suitable for measuring the position and movement of gas in liquid metals [150]. Since radiological methods cannot resolve complex 3D structures, advanced techniques such as X-ray tomography have been developed to retrieve 3D-3C information [123, 84]. However, both X-ray and neutron-based methods are highly specialised and require expensive equipment [150]. Magnetic resonance imaging (MRI) determines velocities by measuring phase shifts in a magnetic field. It enables 3D-3C flow measurements [178, 193, 150] and is particularly useful for flows with opaque continuous phases. However, MRI systems are highly specialised and very expensive. Ultrasonic pulsed Doppler velocimetry (UPDV) [134, 20] measures the frequency shift of moving particles caused by the Doppler effect to determine their velocity. As a line measurement technique (line-integrated velocity profile), it requires movement of the measurement apparatus to cover larger fields. Consequently, UPDV cannot be used to obtain instantaneous volumetric flow fields. Laser Doppler velocimetry (LDV) [220] is a similar technique that uses laser light instead of sound waves. It also relies on the Doppler effect to determine particle velocity [5]. In contrast to UPDV, which provides line-integrated velocity profiles, LDV is a point measurement technique, capturing velocity at a single spatial location where the laser beams intersect. It is therefore limited to stationary or statistically steady flows, where scanning across the field over time does not distort the measurement. An optical measurement technique capable of capturing a field of particles is particle image velocimetry (PIV). PIV uses correlation to track changes in particle patterns between two images, thereby determining the velocity field using an Eulerian approach. As standard PIV can only retrieve a 2D-2C velocity field [4, 156], various developments have extended the method to 3D-3C. Examples of such 3D-3C PIV techniques include tomographic PIV [52] and holographic PIV [59, 32], both of which employ multiple cameras to reconstruct three-dimensional information. For tracking individual particles using a Lagrangian approach, particle tracking velocimetry (PTV) can be employed. In this method, individual particles are detected and their displacement between two images is measured. One class of PTV approaches uses multiple cameras, such as tomographic PTV [143, 126] and subsequent improvements like Shake-the-Box [173, 174]. Alternatively, single-camera methods exploit optical effects, such as defocusing, astigmatism, or numerical refocusing at different depth positions, to reconstruct the third dimension from 2D images. Examples include APTV [33, 86], DPTV [215, 146], and holographic PTV methods [24].

Determining Particle Size

A variety of techniques exist for sizing individual particles. Magnetic resonance imaging (MRI) can be applied not only to determine particle velocities but also to measure particle sizes within three-dimensional volumes [150]. Optical Coherence Tomography (OTC) is another useful technique for measuring both the 3D position and sizes of bubbles [136]. It also offers the advantage of being able to measure more complex bubble or droplet shapes. OTC operates by splitting a broadband light source into a reference beam and a measurement beam, the latter of which is reflected back from the sample. The interference between the reflected measurement beam and the reference beam is then analysed to extract depth-resolved structural information based on optical path length differences. Using this, volumes can be scanned to detect surfaces or interfaces. However, the limited scanning speed of the technique prevents it from being used in large fields of view, although recent improvements in scanning speed could change this in the future. Another technique which can be utilised to measure both the position and size of particles is Light Detection and Ranging (LIDAR) [101]. The method operates by emitting short pulses of laser light towards a target and measuring the time taken for the light to be reflected back to the sensor. Employing a scanning motion similar to that used in OCT, LIDAR is capable of generating three-dimensional maps of particle surfaces, which may be used to reconstruct particle size from a single optical access. However, for sizing small particles (on the order of a few microns), LIDAR relies on indirect measurements based on scattering models, which typically constrain the technique to estimating size distributions rather than resolving individual particles [202]. Imaging methods such as high-speed imaging allow particle sizing within a two-dimensional light sheet. Focused images of particles are captured, and their size is measured directly in pixels and converted using the magnification of the optical system. Since only a thin slice of the flow is in focus, only particles within this plane can be sized. Holography techniques [194, 198] overcome this two-dimensional limitation. These methods, particularly digital holography [102], use digital refocusing at various depth positions to reconstruct 3D volumes. Particle sizes are then extracted from each refocused slice in a manner similar to in-focus imaging. However, both holography and in-focus imaging are limited by camera resolution. Capturing large fields of view requires lower magnification, often without an imaging lens in holography, reducing spatial resolution. This constraint makes it difficult to size small particles in large volumes. Interferometric particle imaging (IPI) addresses this limitation. Originally introduced as interferometric laser imaging for droplet sizing (ILIDS) [98, 64], the method uses interference patterns of scattered light to determine particle diameter from observed fringes. Higher size resolution can be achieved compared to in-focus imaging, but the method is restricted to quasi-spherical particles, where surface tension dominates over volume forces. IPI is generally performed in a two-dimensional light sheet, as the distance between the observer and the particle must be known.

2.2.2 Defocusing Particle Tracking Velocimetry

Particle Tracking

Particle tracking methods can be directly applied to particles of the dispersed phase. However, to obtain the flow field of the continuous phase, the motion of the fluid must first be made visible. This is achieved by introducing tracer particles (or simply tracers) into the fluid - a process known as *seeding*. This concept dates back to Ludwig Prandtl, who used micaceous iron ore particles in water to visualise flow patterns [195]. A seeded continuous phase is, technically, also a two-phase flow. However, since the tracers are solely used to infer the flow field of the carrier phase, they are not considered an additional phase. It is essential to ensure that they accurately represent the fluid

motion. The Stokes number is commonly used to characterise the response time of the tracers, t_P , relative to the response time of the fluid, t_F [197]:

$$\text{Stk} = \frac{t_P}{t_F} = \frac{d_P^2 \rho_P / (18 \mu_F)}{l_c / u_c}. \quad (2.1)$$

To ensure that tracers accurately follow the motion of the carrier phase (i.e. the response time of the particles is sufficiently smaller than the response time of the fluid), a Stokes number smaller than $\text{Stk} < 0.1$ is typically required [156]. The particle response time depends on the particle diameter d_P and density ρ_P , as well as the fluid's dynamic viscosity μ_F . The fluid response time is defined by the characteristic length scale l_c and characteristic velocity u_c of the flow. Seeding is essential for all techniques that rely on tracking particle motion, including Doppler-based methods (UPDV and LDV), PIV, and PTV.

PTV captures images of either tracer particles or particles of the dispersed phase (e.g. droplets). At least two images are recorded at a known time interval. The local velocity of a particle is then directly derived from its displacement between the two images. Setups using more than two images can also provide additional information, such as particle acceleration, which requires evaluating the change in velocity over successive time intervals.

Defocusing Approach

In DPTV, the depth position of a particle is determined from the degree of defocus in its corresponding PI. When a light source, such as light scattered by a particle, is imaged out of focus, it appears as a blurry circle (depending on the aperture shape). This effect is commonly observed in night scenes of films, where background lights appear as blurry disks. The extent of defocusing is related to the distance between the object (i.e. the particle) and the focal plane of the observer (i.e. the camera).

Willert and Gharib [215] first applied this principle to determine a particle's distance from the focal plane by measuring the diameter of its PI, enabling 3D-3C tracking. The relationship between the defocus diameter d_{PI} of the PI and the particle's distance from the focal plane z_P was mathematically described by Olsen and Adrian [146]:

$$d_{PI}^2 = \underbrace{M_G^2 d_P^2}_{\text{geom. image}} + \underbrace{5.95(M_G + 1)^2 \lambda^2 (f^\#)^2}_{\text{diffraction}} + \underbrace{\frac{M_G^2 z_P^2 D_a^2}{(s_0 + z_P)^2}}_{\text{defocusing}}. \quad (2.2)$$

This equation consists of three terms: the first describes the geometrical image, defined by the particle diameter d_P and the system magnification M_G ; the second accounts for diffraction, involving the light wavelength λ and the focal number of the objective lens $f^\#$; and the third represents the defocusing effect, which includes the aperture diameter D_a and the distance from the imaging optics to the focal plane s_0 . In a fixed measurement setup, all terms except z_P remain constant, making d_{PI} a function of the particle's distance from the focal plane: $d_{PI} = f(z_P)$.

In many applications, the imaging optics are positioned much farther from the focal plane than the particles, so that the condition $s_0 \gg z_P$ holds. Under this assumption, the denominator of the defocusing term in Eq. (2.2) can be simplified as $(s_0 + z_P)^2 \approx s_0^2$. The equation then becomes:

$$d_{PI} \sim \left(\text{const.} + M_G^2 D_a^2 z_P^2 \right)^{\frac{1}{2}}. \quad (2.3)$$




Figure 2.3: Defocusing approach visualised. Shown are two particles, one closer to the focal plane (blue) resulting in a small defocus circle and a particle further away from the focal plane (red) resulting in a larger defocus circle. The image plane (i.e. the camera chip) and the (object side) focal plane are shown. Additionally, is the image side focal plane shown, which is the plane at which the defocused particle would be imaged in focus. A 3D close-up on how the defocused particle image changes between the image side focal plane and the actual image plane is shown. Modified figure based on [113].

Eq. (2.2) becomes approximately linear in z_P at sufficient distances from the focal plane z_P . As a result, Eq. (2.3) can be linearised to $d_{PI} \sim z_P$ [57]. This relationship allows the out-of-plane position (i.e. the z -direction) to be directly determined by measuring the PI diameter.

Calibration

In cases where the measurement volume is constrained by two walls, an in-situ calibration can be employed for DPTV [116]. However, this condition is not always met. In such cases, a more general approach using a calibration target can be applied [SDK25]. The relationship between the PI diameter and the out-of-plane position can also be described using geometrical optics (GO) [179]:

$$d_{PI} = D_a \frac{|z_{I2P}^{foc} - z_{I2P}^{defoc}|}{z_{I2P}^{foc}} \quad (2.4)$$

The camera sensor is positioned at a distance z_{I2P}^{foc} from the principal plane, such that it captures a focused image of objects located in the object-side focal plane at a distance z_{P2F} . This configuration defines the nominal imaging geometry, compare the light green rays in Fig. 2.4. When an object is located at a distance z_P from this object-side focal plane, its image is no longer formed at the camera sensor but instead at a different location, denoted as z_{I2P}^{defoc} , compare the dark green rays in Fig. 2.4. As a result, the image recorded on the sensor appears defocused. Eq. (2.4) shows good agreement with the linear part of Eq. (2.2), as will be shown later (e.g. Fig. 4.4) [SGK25]. However, the position of the principal plane is typically unknown. Therefore, both z_{I2P} and z_{P2F} must first be determined. To this end, the thin lens equation

$$\frac{1}{z_{I2P}^{foc}} + \frac{1}{z_{P2F}} = \frac{1}{f_{lens}} \quad (2.5)$$




Figure 2.4: "Working Principle of DPTV visualised. As the particle moves away from the focal plane the particle image visible on the image plane becomes defocused and the diameter of the PI increases." [SGK25]. Figure from [SGK25]

and the magnification condition

$$M_G = \frac{|z_{I2P}^{foc}|}{z_{P2F}} \quad (2.6)$$

can be utilised. While the focal length of the imaging lens is usually known, the magnification can be determined from a calibration target. Once z_{I2P}^{foc} is known, z_{I2P}^{defoc} can be obtained from Eq. (2.4). Substituting z_{I2P}^{defoc} into the thin lens equation allows the derivation of the distance from the principal plane to the particle, i.e., $z_{P2F} + z_P$. Knowing both this total distance and the distance from the focal plane to the principal plane, (z_{P2F}), the particle's distance from the focal plane, (z_P), is simply the difference between the two. This procedure is employed in Chapter 3 and Chapter 4.

Particle Image Density

When evaluating PTV methods, the number of particles N_P in the image plays a crucial role. In PIV and PTV, this is commonly referred to as the seeding density [3]. In PTV, the seeding density is particularly important because it determines the spatial resolution of the flow field by setting the number of velocity vectors in the image (see Fig. 2.5). Specifically, the number of individual velocity vectors per volume is directly influenced by the number of particles present within that volume.

During image post-processing, individual particles must be detected and matched between images. Both detection and tracking algorithms have limitations regarding the number of particles they can handle before performance is impaired, leading to missed detections and/or incorrect matches. These issues can result in ghost particles or erroneous velocity vectors, respectively. Therefore, the seeding density of tracers in the continuous phase should be adjusted to balance spatial resolution with algorithmic limitations. In DPTV, an additional factor must be considered: since PIs are defocused, they appear as relatively large circles compared to the small points seen in in-focus methods, compare Fig. 2.5. This can lead to *particle image overlap*, where the PIs of two or more particles intersect, complicating particle detection. To address this, a more suitable metric for




Figure 2.5: Resolution of a flow field, based on the number of particles in the fluid visualised. A flow field of a vortex is shown as an example. The left image contains $N_p = 20$ particles and the right image contains 60 particles. The effect of defocusing and the seeding density on the particle image overlap is shown.

seeding density is required in DPTV. Kähler *et al.* [105] proposed relating the total area of the PIs to the full image area. The particle image density is defined as [12]:

$$\mathcal{S} = \frac{1}{A_{\text{image}}} \sum_{\ell=1}^{N_P} A_{\text{PI},\ell} \approx N_P \frac{\overline{A_{\text{PI}}}}{A_{\text{image}}} \quad (2.7)$$

where A_{image} is the area of the full image, $A_{\text{PI},\ell}$ is the area of the ℓ -th PI, and $\overline{A_{\text{PI}}}$ is the mean area of all PIs. The problem of PI overlap and seeding density is addressed in more detail in Chapter 6.

Post-Processing

To obtain particle velocities in DPTV, three main steps must be followed. First, the PI must be detected (detection step), which involves determining the in-plane position $(x_{\text{PI}}, y_{\text{PI}})$ and the diameter d_{PI} of the PI. In the second step, the physical 3D position of the particle (x_P, y_P, z_P) is reconstructed from the PI. Once the physical position is known, the final step is to match the particle between two images to determine its velocity (tracking step).

There are many different approaches to particle detection. A brief overview is provided below: "One straight-forward approach used by Leister *et al.* [114, 115] is the Coherent Hough Transform [...] [7, 9, 8] – a variation of the Hough Transform [79] – which is a gradient based voting algorithm. Fuchs *et al.* [57] used an algorithm detecting the edges of the defocus rings by an adaptive threshold applied on the intensity distribution. While the first two methods determine the defocus rings diameter and center position directly from the particle image, Barnkob *et al.* [13, 12, 165] used cross correlation to compare the measured particle images to reference ones from a calibration stack, both for particle detection and z -position refinement." [SDLK23]. "With the rise of machine learning in computer vision and neural networks in particular, new methods for particle detection based on convolutional neural networks (CNN) [109, 111] emerged in the DPTV/APTV community. Cierpka *et al.* [35] demonstrated the applicability of Faster R-CNN [164] for particle detection in APTV. Franchini *et al.* [56] demonstrated an improved detection rate on overlapping

particle images for a CNN-based model in APTV. König *et al.* [99] used a cascaded CNN on the basis of Faster R-CNN on APTV, which showed to have lower position uncertainties on particles with astigmatisms and noise compared to conventional algorithms. Barnkob *et al.* [14] also used a cascaded version of a CNN for particle detection in DPTV and APTV, which was composed of a preliminary network (Faster R-CNN) for locating particles within the image and a second CNN to refine the out-of-plane position. More recently, Dreisbach *et al.* [48] applied CNN-based multi-stage (Faster R-CNN) and single-stage detectors (RetinaNet [119]) for DPTV, and furthermore analyzed the effect of synthetic training data refinement by means of generative adversarial networks (GAN) on the network performance for DPTV image processing."[SDLK23] Finally, Sax *et al.* [SDLK23] proposed a hybrid method in which the initial detection by a neural network is refined using edge detection, and a physical check criterion is applied to eliminate incorrect detections. Since these initial efforts to incorporate machine learning into DPTV, CNNs have increasingly gained traction within the DPTV community, see e.g. [223, 159].

The second step involves reconstructing the physical position of the particle from the PI. The following relationships apply:

$$\begin{aligned} \text{in-plane position} \quad (x_{\text{PI}}, y_{\text{PI}}) &= (x_{\text{P}}, y_{\text{P}})M_G \\ \text{out-of-plane position} \quad z_{\text{P}} &= f(d_{\text{PI}}) \end{aligned} \tag{2.8}$$

The in-plane position is reconstructed directly from the centre coordinates of the PI, whereas the out-of-plane position is derived from the PI diameter, as previously described.

Various approaches exist for 3D particle tracking. The simplest is the nearest neighbour method, which matches particles between images based on proximity. This approach requires that particle displacements between frames are smaller than the distance to neighbouring particles to ensure unambiguous matching. More advanced methods use multiple images and incorporate velocity vector direction to optimise matching in subsequent time steps. An example of such a method is Shake-the-Box [173].

Limitations in DPTV

The post-processing of images in DPTV is subject to several key limitations, both optical and algorithmic in nature. First, PI overlap can significantly impair detection performance. Overlapping PIs reduce the number of correctly identified particles (i.e., lower recall) and increase the likelihood of false positives (i.e., lower precision). An inherent trade-off exists: larger particles are generally better spatially resolved and, in principle, allow for more accurate detection and size estimation; however, they also elevate the risk of overlap, particularly at high particle image densities or under increased defocusing. This issue is especially pronounced in setups with high defocus sensitivity (i.e., a large change in PI size with respect to z_{P}). Second, the signal-to-noise ratio (SNR) plays a critical role in detection accuracy. As defocus increases, the light from a particle is distributed over a larger area, reducing the peak intensity and rendering the PI less distinguishable from the background. This degradation in contrast can hinder both reliable detection and accurate PI size determination. These limitations must be carefully considered when applying DPTV, as they can lead to detection errors and, consequently, inaccurate velocity vector fields.

2.2.3 Interferometric Particle Imaging

The next technique, which is a central focus of Chapters 3, 4 and 5, is IPI. IPI incorporates several optical concepts, such as the scattering of light by transparent particles, which is discussed in greater detail in Section 2.3, interference and diffraction. However, a brief overview is provided here to offer the necessary background for understanding the working principle of IPI. IPI addresses the limitation of size resolution in in-focus imaging techniques when applied to comparatively large fields of view.

Originally introduced as ILIDS [98, 64, 139], the technique was developed specifically for droplet sizing. It utilises the interference of coherent monochromatic light scattered by a quasi-spherical transparent particle. When this particle is observed in an out-of-focus plane, a PI similar to those in DPTV can be seen. The key difference is that, under certain observation angles, the PI contains an interference pattern. This pattern provides information about the particle size, analogous to how fringes relate to slit spacing in a Young's fringe experiment. The method has been widely used for droplet characterisation. "Examples of the use of IPI on droplet flows include the characterization of sprays [128, 70], droplets analysis at nozzle outlets [132], droplets in combustion [58] and droplets analysis in clouds and the atmosphere [50, 93, 154]"[SDK25]. The technique was later extended from droplets to bubbles and renamed IPI [90, 144]. Since then, IPI has also been employed for sizing small bubbles [144, 90, 106]. "Further developments include the consideration of higher-order frequencies in the interference patterns to create redundancy [42] and therefore increased precision, and the generalization of the IPI formula to arbitrary paraxial optical systems [179]."[SDK25].

Geometrical Optics Model for Light Scattering at a transparent Particle

In order to understand the working principle of IPI, two basic concepts must first be introduced: light scattering at particles and the concept of glare points. Both are discussed in greater detail in Section 2.3.1.

"First, only the scattering process at a bubble or droplet is regarded. Let's consider a single transparent, approximately spherical particle with a refractive index n_1 that is different from the refractive index of the surrounding n_2 . The ratio between the refractive indices is then described by $m = n_1/n_2$ (i.e. relative refractive index). The scattering of light at the particle is initially described by the geometrical optics (GO) model. Light rays are either reflected or refracted at the particle surface, compare Fig. 2.6. A light ray directly reflected at the particle surface is denoted as order zero $p = 0$, while refracted light entering the particle is of order $p \geq 1$. The scattering order of the refracted light describes the number of chords traveled inside the particle or $p - 1$ internal reflections. The exit points of the light rays are called glare points and are perceived as bright spots on the particle surface when viewed from a scattering angle θ [200, 201], see Fig. 2.7."[SDK25]

A laser light beam, consisting of many parallel rays, will be scattered in all directions, resulting in glare points that are visible from almost every angle. "Only light rays that exit the particle in parallel with the angle θ eventually reach the camera and thus all other rays can be discarded for this model. Furthermore, refraction at the particle surface is governed by Snell's law, which leads to two fundamental implications for the scattering process: The positions of the glare points depend on the scattering angle, see Fig. 2.6, and on the relative refractive index, compare Fig. 2.7. This mechanism causes in particular a disparity between particles with $m > 1$ (droplets) and $m < 1$ (bubbles). Furthermore, the intensity of the glare points also depends on the scattering angle, due to the angle-dependent transmission and reflection coefficients at the particle surface, as described




Figure 2.6: Light scattering on a spherical particle as visualised by the geometrical optics model. The relationship between glare points and the scattering angle θ is illustrated. The scattering of the dominant light rays is indicated on the blue spheres. The positions of the glare points, as they would appear to an observer at the angle θ , can be seen on the black spheres. Figure adapted from [SDK25].




Figure 2.7: "A bubble $m < 1$ (left) and a droplet $m > 1$ (right) are shown. Depicted are the reflected ray $p = 0$ (●) and the first three orders of refracted rays $p = 1$ (●), $p = 2$ (●), and $p = 3$ (●). The position of the glare point on the sphere's surface is indicated by the w coordinate. Note that the light is monochromatic and the colours encode the scattering order." [SDK25]. Figure adapted from [SDK25]

by Fresnel's formulae. Therefore, different scattering angles are dominated by different scattering orders (usually not more than one or two orders at a time). [SDK25]

"In a plane scattering process, glare points are located on the equator of the sphere. In analogy to the work of Van de Hulst *et al.* [200, 201], the position of a glare point on the particle surface is denoted by the normalized w -coordinate, which represents the projection of the sphere's equator onto a straight line as seen by a camera, compare Fig. 2.6. The position w_i of a glare point of order p_i , can take values of $[-1, 1]$, with $|w| = 0$ representing the center and $|w| = 1$ the edges of the sphere as seen by the camera. With the GO model, w_i can be determined by geometrical considerations, as is illustrated in Fig. 2.8a. Based on the commonly used convention [98, 90], the angle between the light ray and the tangent of the particle surface is defined as β . The angle of an incident ray of order p_i is denoted as $\beta_i^{(p_i)}$ and the angle of the transmitted ray is denoted as $\beta_t^{(p_i)}$. In the case of the reflected ($p = 0$) light ray, β_r is used for the incident ray. As the scattering angle




Figure 2.8: "(a) Light paths of the $p = 0$ and $p = 1$ rays through the particle using the GO model to determine the position of the glare point, on the example of a droplet. The working principle of IPI is visualized in (b). From the observation angle θ , two glare points are visible on the particle surface, when imaged in focus. The images of the glare points begin to overlap and form an interference pattern as they move away from the focal plane. They eventually collapse into a single circle at a sufficient distance from the focal plane. In addition, the analogy of the glare point spacing (particle size) in IPI to a double slit in Young's fringe experiment [73] is visualized ((b) top). The intensity maxima are depicted as dark stripes for better visibility (inverted intensity)."[SDK25]. Figure adapted from [SDK25]."

θ is assumed to be known, Snell's law together with the sum of the angles provides two equations for the calculation of the two angles β_i and β_t . The glare point position is then calculated from β_i and the particle diameter

$$w_{p_i}^{\text{phys}} = \frac{d_P}{2} w^{(p_i)} = \frac{d_P}{2} q \cos(\beta_i^{(p_i)}), \quad (2.9)$$

where $w_{p_i}^{\text{phys}}$ is the distance in physical units (e.g. μm), $w^{(p_i)}$ is the normalized position, and a factor $q = \pm 1$ indicates on which side of the sphere the glare point is located."[SDK25]

Transfer Functions in IPI

"When the particle is illuminated with coherent, monochromatic light, the light scattered off the particle forms an interference pattern in the far field (distance from the particle $\gg d_p^2/\lambda$, with d_p being the particle size [121]). The underlying physical process is comparable to the one observed at a double slit in Young's fringe experiment, see Fig. 2.8b. While the scattered light is visible as glare points when imaged in-focus, the glare points become defocused rings when the observer is not located on the focal plane. At a sufficient distance from the focal plane, the defocus rings overlap to form the interference pattern, see Fig. 2.8b (bottom).

The particle's diameter can then be reconstructed from the relationship between the number of stripes N_S in the interference pattern and the distance between the glare points. To relate the glare point spacing to the number of stripes, either the distance between the particle and the observer or the opening angle of the aperture α_{AP} must be known. The relationship between the particle diameter and the number of stripes can be described for droplets ($m > 1$) [167, 139] by

$$d_p = \frac{2\lambda N_S}{\alpha_{AP}} \left(\cos(\theta/2) + \frac{m \sin(\theta/2)}{\sqrt{m^2 - 2m \cos(\theta/2) + 1}} \right)^{-1} \quad (2.10)$$

and for bubbles ($m < 1$) [144, 90] by [SDK25]

$$d_p = \frac{2\lambda N_S}{\alpha_{AP}} \left(\cos(\theta/2) - \frac{m \sin(\theta/2)}{\sqrt{m^2 - 2m \cos(\theta/2) + 1}} \right)^{-1}. \quad (2.11)$$

Eqs. (2.10) and (2.11) are frequently employed to determine the diameter of a droplet or bubble based on the number of fringes observed in the PI. An alternative to directly counting the fringes in the PI involves deriving the fringe count indirectly from the fringe frequency, utilising fast Fourier transforms (FFT) and techniques such as peak-fitting in the frequency domain. The latter approaches have been demonstrated to offer greater precision than direct fringe counting [90, 144, 153].

However, Eqs. (2.10) and (2.11) are applicable solely within the front-scatter regime, as discussed in Chapter 3.

2.3 Relevant Concepts in Optics

This chapter introduces theoretical concepts in optics that are employed throughout this work. First, the scattering of light waves by transparent particles such as droplets and bubbles is described in Section 2.3.1. These concepts are used in the Chapters 3 and 4. Second, the ray transfer matrix (RTM) model is introduced in Section 2.3.2. The RTM provides a geometrical description of light propagation in paraxial optical systems, such as those encountered in DPTV and IPI systems. Finally, Section 2.3.3 presents fundamental concepts related to the propagation of light waves, which are relevant to model DPTV and IPI measurements in Chapter 4.

2.3.1 Light Scattering at Droplets and Bubbles

This section reviews light scattering concepts relevant to bubbles and droplets of sizes relevant to DTPFs (i.e. significantly larger than molecules). The content of this section is applied in Chapter 3 to develop a model for IPI in backscatter, and in Chapter 4 to construct the forward model for simulating particle images.

There are numerous scattering theories, each with its own range of validity. Scattering at molecules is described by Raman scattering [157]. Scattering at particles smaller than approximately a tenth of the wavelength of light can be described by Rayleigh scattering [160]. Particles with sizes comparable to the wavelength of light can be described using Mie scattering [137, 200]. The Lorenz-Mie theory has no upper limit to its validity and is therefore often used to describe the scattering of light by small droplets and bubbles. However, computing the Mie solution for large particles becomes computationally expensive [216], making it impractical for very large particles. For such large particles, the geometrical optics approach can be employed to describe the scattering process without significant approximation errors [201].

This section begins by introducing light scattering at spherical homogeneous spheres, using the geometrical optics model. It then presents the Lorenz-Mie theory to describe the scattering process in terms of wave optics. The third subsection introduces the Debye series expansion of the Mie theory to investigate individual scattering orders. The final subsection discusses glare points in the context of wave scattering at spherical bubbles or droplets.

Geometrical Optics Approach for Light Scattering

The concept of describing light scattering at particles using the GO model was already introduced in Section 2.2.3. This chapter expands on the theory and introduces additional concepts used throughout this thesis. Both the GO and wave optics models are employed in this work. The GO scattering model is particularly useful, as it provides a more intuitive representation of certain scattering phenomena than the wave optics model.

Using Snell's law [19, 124], the transmission angle can be expressed in terms of the angle of incidence and the relative refractive index:

$$\beta_t = \arccos\left(\frac{1}{m} \cos(\beta_i)\right) \quad (2.12)$$

Van de Hulst [201] introduced an equation for the incident angle of a light ray at a sphere in relation to the scattering angle θ and the total deflection angle θ_{tot} :

$$\theta_{\text{tot}} = 2 \left(\beta_i^{(p)} - p \arccos\left(\frac{1}{m} \cos(\beta_i^{(p)})\right) \right) = 2\pi k + q\theta \quad (2.13)$$

where $q = \pm 1$ and $k \in \mathbb{N}$ is a constant. This relationship allows the prediction of glare point positions for different scattering orders p . Notably, Eq. (2.13) has more than one solution for certain scattering orders and refractive indices, meaning that more than one glare point of the same order can exist. A particularly interesting phenomenon occurs when multiple rays of the same order collapse into a single glare point at specific scattering angles. These angles are referred to as rainbow angles. Rainbow angles occur under the condition $d\theta/dw = 0$ and are typically observed for droplets and for $p = 2, 3$ [201]. Glare points at rainbow angles are notable for their significantly increased intensity and are the reason why rainbows can be seen in the sky [201].

So far, only rays being either reflected or refracted have been considered. However, at each interaction of a ray with an interface, the ray is both partially reflected and transmitted. A ray interacting with a sphere is shown in Fig. 2.9. When a plane wave encounters a boundary between two homogeneous media, it splits into two components: a transmitted ray that continues into the second medium, and a reflected ray that propagates back into the original medium [19]. The amplitudes of the reflected and transmitted waves can be described by the reflection coefficient $R_{1,2}$ and the transmission coefficient $T_{1,2}$, where the subscript 1 denotes the transverse electric (TE) mode (also called perpendicular polarisation), and 2 denotes the transverse magnetic (TM) mode (also called parallel polarisation). The reflection and transmission coefficients are given by the Fresnel equations [124, 19], which are stated below for non-magnetic media:

$$\begin{aligned} R_1 &= -\frac{\sin(\beta_i - \beta_t)}{\sin(\beta_i + \beta_t)} \\ R_2 &= -\frac{\tan(\beta_i - \beta_t)}{\tan(\beta_i + \beta_t)} \end{aligned} \quad (2.14)$$

$$\begin{aligned} T_1 &= -\frac{2 \sin(\beta_t) \cos(\beta_i)}{\sin(\beta_i + \beta_t)} = 1 - R_1 \\ T_2 &= -\frac{2 \sin(\beta_t) \cos(\beta_i)}{\sin(\beta_i + \beta_t) \cos(\beta_i - \beta_t)} = 1 - R_2 \end{aligned} \quad (2.15)$$

where β_i is the incident angle of the light and β_t is the angle of the transmitted light.




Figure 2.9: A light ray (scattering order $p = 1$) interacting with a sphere of refractive index n_1 in a surrounding of refractive index n_2 (left). At every interaction of the ray with the interface a part of the ray is transmitted and a part reflected. The transmission and reflection coefficients at each interaction with an interface are shown, with e.g. T^{12} denoting a transmission from medium one (n_1) to medium two (n_2) and R^{11} denoting a reflection from medium one back to medium one. The different light intensity that is observed based on the scattering angle is visualised in the second figure (right).

For parallel polarised light, there exists an incident angle at which the transmission coefficient becomes one, $T_2 = 1$, and all light is transmitted. This angle is known as the Brewster angle. Using Eq. (2.15) and Snell's law, Brewster's law is derived [124]:

$$\beta_{i,B} = \arctan\left(\frac{n_2}{n_1}\right) \quad (2.16)$$

where $\beta_{i,B}$ is the Brewster angle, n_1 is the refractive index of the original medium (medium of incidence), and n_2 is the refractive index of the medium into which the light is transmitted.

When a light ray from the surrounding medium hits the interface of the sphere, the transmission coefficient from medium n_2 into n_1 (into the sphere) is denoted by T^{21} . Note that the superscript does not indicate a power, but rather denotes the transition from medium two to medium one. A part of the ray is reflected back into medium two with coefficient R^{22} , see the left side of Fig. 2.9. The transmitted ray is refracted into the sphere and continues to travel until it interacts again with the interface. A part T^{12} of the ray is then transmitted out of the sphere (from medium one to two). This ray, upon exiting the sphere, has an intensity of $T^{21}T^{12}$ relative to the original incoming ray. The remaining part of the ray is internally reflected with coefficient R^{11} and has an intensity of $T^{21}R^{11}$ with respect to the original ray. The internally reflected ray continues to propagate until it encounters the next interface, and the process continues. Since the reflection and transmission coefficients are functions of the angle of incidence relative to the sphere's surface, the intensity of the emergent ray exhibits a strong dependence on the scattering angle θ (see the right panel of Fig. 2.9). Owing to partial transmission and reflection at each interface interaction, the ray undergoes a progressive loss of intensity with every internal reflection. Consequently, rays of higher scattering order, those undergoing multiple internal reflections, contribute progressively less to the overall scattered intensity and very large scattering orders can eventually be considered negligible.

The Lorenz-Mie Theory

The previous section introduced light scattering at spheres using the GO model. While the GO model is valuable for determining the positions of glare points and for developing an intuitive understanding of scattering phenomena, it does not fully account for all physical effects required to accurately compute the intensity of scattered light. To address these limitations, this section introduces light scattering at spheres from the perspective of wave optics. The Lorenz-Mie theory

(LMT) provides a widely used exact solution for the scattering of plane electromagnetic waves by homogeneous spheres. It is derived from Maxwell's equations formulated in polar coordinates [200]. The LMT yields solutions for both the near-field and far-field scattering regimes. However, since an observer relevant for IPI is usually located sufficiently far from the particle (i.e. at a distance several orders of magnitude greater than the wavelength of light), only the solution for the far field is considered here. The traditional LMT describes the integral light intensity over all scattering orders.

"The LMT uses the previously introduced relative refractive index m , where $\Re\{m\} = n_1/n_2$ and $\Im\{m\}$ describes the attenuation of the wave for partially opaque materials. In addition, a size factor $x_{\text{size}} = \pi d_p/(\lambda_0/n_2)$ is introduced for the dimensionless parametrization of the particle radius." [SDK25] The size factor relates the particle diameter d_p to the wavelength of the incident wave in the surrounding medium, λ_0/n_2 , where λ_0 denotes the wavelength in vacuum and n_2 represents the refractive index of the medium surrounding the sphere. The complex amplitude of the scattered light in the far field, $|S| = I$, was derived by van de Hulst [200] from the LMT. The intensity I of the scattered light is composed of the contributions from the TE mode, denoted by (S_1) and the TM mode, denoted by (S_2).

$$I = \sqrt{S_1^2(\theta) + S_2^2(\theta)}. \quad (2.17)$$

The solution of the LMT in the far field for a homogeneous sphere of size x_{size} and relative refractive index m looks as follows [200]:

$$S_1(\theta) = \sum_{\ell=1}^{\infty} \frac{2\ell+1}{\ell(\ell+1)} (a_{\ell}\Pi_{\ell}(\theta) + b_{\ell}\tau_{\ell}(\theta)) \quad (2.18)$$

$$S_2(\theta) = \sum_{\ell=1}^{\infty} \frac{2\ell+1}{\ell(\ell+1)} (b_{\ell}\Pi_{\ell}(\theta) + a_{\ell}\tau_{\ell}(\theta)) \quad (2.19)$$

where a_{ℓ} and b_{ℓ} are the Mie coefficients which are functions of the relative refractive index m and the size of the sphere x_{size} . In the LMT, the incident and the scattered wave are described by a series of spherical wave functions (i.e. partial waves or multipole modes ℓ). The functions $\Pi_{\ell}(\theta)$ and $\tau_{\ell}(\theta)$ are known as the Mie angular functions; they describe the dependence of the scattered field on the scattering angle and characterise its angular distribution. Moreover, they represent the angular structure associated with each multipole mode indexed by ℓ . The Mie coefficients can be interpreted as weighting factors that quantify the contribution of the ℓ^{th} Mie angular function for a given m and x_{size} . Physically, a_{ℓ} and b_{ℓ} represent the excitation amplitudes of the electric (TM) and magnetic (TE) multipole modes, respectively. The Mie angular functions $\Pi_{\ell}(\theta)$ and $\tau_{\ell}(\theta)$ are visualised for the orders $\ell = [2, 6]$ in Fig. 2.10.

The number of partial waves that should be considered to compute the Mie solution depends on the size of the particle [217, 145, 18]:

$$\ell_{\text{stop}} = x_{\text{size}} + 4.05x_{\text{size}}^{1/3} + 2 \quad (2.20)$$

For particles relevant to IPI, the following assumptions can be made: First, the sphere is large compared to the wavelength of the light ($x_{\text{size}} \gg 1$), therefore, a large number of partial waves are considered ($\ell \gg 1$). Second, the boundary cases $\theta \neq \{0, 180\}$ are excluded. Then, $\Pi_{\ell}(\theta) \ll \tau_{\ell}(\theta)$ applies, and therefore S_1 becomes associated with the TE-mode and S_2 with the TM-mode [200] (perpendicular and parallel polarisation, respectively). The full solution by the LMT for an air




Figure 2.10: Mie angular functions $\Pi_\ell(\theta)$ (top) and $\tau_\ell(\theta)$ (bottom) visualised for the partial waves of order $\ell = 2, \dots, 6$.




Figure 2.11: Logarithmic intensity of light scattered by an air bubble in water over the scattering angle θ , according to the Mie solution. The depicted bubble has a diameter $d_P = 20 \mu\text{m}$, $n_1 = 1$, $n_2 = 1.333$ and the light has a wavelength of $\lambda = 532 \text{ nm}$, ($m = 1/1.333$, $x_{\text{size}} = 314.9$).

bubble in water, which is a weighted superposition of spherical functions like in Fig. 2.10 is shown in Fig. 2.11.

The Mie angular functions for a given order ℓ can be computed from the associated Legendre functions of the first order and degree ℓ [200].

$$\Pi_\ell(\theta) = \frac{P_\ell^{(1)}(\cos(\theta))}{\sin(\theta)} \quad (2.21)$$

$$\tau_\ell(\theta) = \frac{d}{d\theta} P_\ell^{(1)}(\cos(\theta)) \quad (2.22)$$

The associated Legendre polynomials can be expressed by Rodrigue's formula[2]:

$$P_\ell^{(1)}(z) = \frac{1}{2^\ell \ell!} \frac{d}{dz} (z^2 - 1)^\ell \quad (2.23)$$

Using identity 8.5.2 from Abramowitz and Stegun [2], an expression for the derivative of the associated Legendre polynomial can also be written in terms of the zero and first order associated Legendre polynomials:

$$\frac{dP_\ell^{(1)}(z)}{dz} = \frac{\ell(\ell+1)}{(z^2-1)^{1/2}}P_\ell^{(0)}(z) - \frac{z}{z^2-1}P_\ell^{(1)}(z) \quad (2.24)$$

More information on the associated Legendre functions can be found in Abramowitz and Stegun [2].

While the Mie angular functions describe the angular dependency of the scattering process, the Mie coefficients encompass the dependence on the particle size and relative refractive index. The Mie coefficients a_ℓ and b_ℓ can be computed from the spherical Bessel functions and spherical Hankel functions [133]:

$$a_\ell = \frac{m^2 \mathcal{J}_\ell(mx_{\text{size}})[x_{\text{size}}\mathcal{J}_\ell(x_{\text{size}})]' - \mathcal{J}_\ell(x_{\text{size}})[mx_{\text{size}}\mathcal{J}_\ell(mx_{\text{size}})]'}{m^2 \mathcal{J}_\ell(mx_{\text{size}})[x_{\text{size}}\mathcal{H}_\ell^{(1)}(x_{\text{size}})]' - \mathcal{H}_\ell^{(1)}(x_{\text{size}})[mx_{\text{size}}\mathcal{J}_\ell(mx_{\text{size}})]'} \quad (2.25)$$

$$b_\ell = \frac{\mathcal{J}_\ell(mx_{\text{size}})[x_{\text{size}}\mathcal{J}_\ell(x_{\text{size}})]' - \mathcal{J}_\ell(x_{\text{size}})[mx_{\text{size}}\mathcal{J}_\ell(mx_{\text{size}})]'}{\mathcal{J}_\ell(mx_{\text{size}})[x_{\text{size}}\mathcal{H}_\ell^{(1)}(x_{\text{size}})]' - \mathcal{H}_\ell^{(1)}(x_{\text{size}})[mx_{\text{size}}\mathcal{J}_\ell(mx_{\text{size}})]'} \quad (2.26)$$

where $\mathcal{J}_\ell(z)$ is the spherical Bessel function of the first kind and order ℓ , $\mathcal{H}_\ell^{(1)}(z)$ is the spherical Hankel function of the first kind, and the prime denotes the derivative with respect to the argument (i.e. x_{size} or mx_{size}) of the Bessel function. The derivatives can be expressed as follows using a characteristic property of the Bessel functions [133]:

$$[z\mathcal{J}_\ell(z)]' = z\mathcal{J}_{\ell-1}(z) - \ell\mathcal{J}_\ell(z) \quad (2.27)$$

and

$$[z\mathcal{H}_\ell^{(1)}(z)]' = z\mathcal{H}_{\ell-1}^{(1)}(z) - \ell\mathcal{H}_\ell^{(1)}(z). \quad (2.28)$$

The spherical Hankel function of the first kind is computed from the spherical Bessel functions of the first and second kind, $\mathcal{Y}_\ell(z)$, with $\mathcal{H}_\ell^{(1)}(z) = \mathcal{J}_\ell(z) + i\mathcal{Y}_\ell(z)$.

The Debye Series Expansion of the Lorenz-Mie-Theory

The LMT only describes the integral intensity over all scattering orders. In order to obtain information on the individual scattering orders p , the Debye Series Expansion (DSE)[40] of the LMT expands the Mie coefficients. The structure of the series expansion of the Mie coefficients follows the idea from the transmission and reflection coefficients introduced in the GO model of scattering at a sphere. The Mie coefficients in the DSE can be expressed in the following form [68, 80]:

$$\begin{bmatrix} a_\ell \\ b_\ell \end{bmatrix} = \frac{1}{2} \left(1 - R_{\ell,a,b}^{22} - \sum_{p=1}^{\infty} T_{\ell,a,b}^{21} (R_{\ell,a,b}^{11})^{p-1} T_{\ell,a,b}^{12} \right) \quad (2.29)$$

The transmission and reflection coefficients have the same physical interpretation as in the GO model.

With $T_\ell = T_\ell^{12}T_\ell^{21}$, the Mie coefficients for scattering order p can be written as

$$\begin{aligned} a_\ell^{(p=0)} &= \frac{1}{2}(1 - R_{\ell,a}^{22}) \\ b_\ell^{(p=0)} &= \frac{1}{2}(1 - R_{\ell,b}^{22}) \end{aligned} \quad (2.30)$$

$$\begin{aligned} a_\ell^{(p>0)} &= -\frac{1}{2}T_{\ell,a}^{21}(R_{\ell,a}^{11})^{p-1}T_{\ell,a}^{21} = -\frac{1}{2}T_{\ell,a}(R_{\ell,a}^{11})^{p-1} \\ b_\ell^{(p>0)} &= -\frac{1}{2}T_{\ell,b}^{21}(R_{\ell,b}^{11})^{p-1}T_{\ell,b}^{21} = -\frac{1}{2}T_{\ell,b}(R_{\ell,b}^{11})^{p-1} \end{aligned} \quad (2.31)$$

since every wave that is transmitted into the sphere must eventually also be transmitted out of the sphere in order to reach the observer. The scattering intensity of light scattered at a bubble, decomposed into the scattering orders $p = [0, 1, 2]$ in TM-mode, is shown in Fig. 2.12.




Figure 2.12: Logarithmic intensity of light scattered by an air bubble in water over the scattering angle θ . The intensity of the first three scattering orders ($p = 0, 1, 2$) in TM-mode, according to the DSE of the LMT are shown. All intensities are scaled identically to visualise the intensity loss of higher scattering orders. The depicted bubble has a diameter $d_P = 20 \mu\text{m}$, $n_1 = 1$, $n_2 = 1.333$ and the light has a wavelength of $\lambda = 532 \text{ nm}$, ($m = 1/1.333$, $x_{\text{size}} = 314.9$).

The reflection coefficients R and transmission coefficients T for incident waves can be calculated from the Riccati-Bessel functions of the third kind $\Gamma_\ell(z)$ and $\Upsilon_\ell(z)$ respectively, and their derivatives with respect to the argument x_{size} or $m x_{\text{size}}$, $\Gamma'_\ell(z)$ and $\Upsilon'_\ell(z)$, as shown in Eqs. (2.32) and (2.33) [68]:

$$R_{\ell a,b}^{22} = \frac{1}{D_{a,b}}[A\Gamma'_\ell(x_{\text{size}})\Gamma_\ell(m x_{\text{size}}) - B\Gamma_\ell(x_{\text{size}})\Gamma'_\ell(m x_{\text{size}})] \quad (2.32)$$

$$T_{\ell a,b}^{21} = 2i \frac{B}{D_{a,b}} \quad (2.33)$$

The abbreviation $D_{a,b}$ can be introduced as shown in Equation 2.34.

$$D_{a,b} = -A\Upsilon'_\ell(x_{\text{size}})\Gamma_\ell(m x_{\text{size}}) + B\Upsilon_\ell(x_{\text{size}})\Gamma'_\ell(m x_{\text{size}}) \quad (2.34)$$

The following abbreviations are used for better readability:

$$\begin{aligned} A &= m; B = 1 \text{ for TM - wave } a_\ell \\ A &= 1; B = m \text{ for TE - wave } b_\ell \end{aligned} \quad (2.35)$$

For out-going waves the calculation follows accordingly with Eqs. 2.36 and 2.33 [68].

$$R_{\ell a,b}^{11} = \frac{1}{D_{a,b}}[A\Upsilon'_\ell(x_{\text{size}})\Upsilon_\ell(m x_{\text{size}}) - B\Upsilon_\ell(x_{\text{size}})\Upsilon'_\ell(m x_{\text{size}})] \quad (2.36)$$

$$T_{\ell a,b}^{12} = 2i \frac{A}{D_{a,b}} \quad (2.37)$$

The Riccati-Bessel functions of the third kind $\Gamma_\ell(z)$ and $\Upsilon_\ell(z)$ can be computed from the spherical Hankel functions[2]

$$\begin{pmatrix} \Gamma_\ell(z) \\ \Upsilon_\ell(z) \end{pmatrix} = \begin{pmatrix} z\mathcal{H}_\ell^{(1)}(z) \\ z\mathcal{H}_\ell^{(2)}(z) \end{pmatrix} \quad (2.38)$$

More information on the efficient computation of the Mie angular functions and Mie coefficients for the DSE of the LMT using stable recursion can be found in Shen *et al.* [182] and Lentz [117]. By using Eq. (2.29) for the Mie coefficients instead of Eq. (2.25) and (2.26) in Eq. (2.18) and (2.19), the scattering intensity of individual scattering orders can be computed instead of the integral Mie solution.

Glare Points in the Lorenz-Mie Theory

The concept of glare points was introduced in Section 2.2.3 using the GO model. The GO model offers a simplified, yet often sufficient, framework for describing glare points. However, a more comprehensive explanation of the glare point phenomenon is provided by the Lorenz-Mie Theory (LMT).

Van de Hulst and Wang [201] proposed a method for computing both the position and intensity of glare points using the Fourier optics approximation. In their formulation, a thin lens is assumed to be positioned in the far-field of the scattered wave. The scattered light propagating along the optical axis is then focused by this lens and projected onto an observation plane situated along the same axis, behind the lens. The complex field $A_{\text{GP},1,2}(w)$ of the scattered wave in this observation plane is expressed as

$$A_{\text{GP},1,2}(w) = \int_{\theta_0 - \Delta\theta}^{\theta_0 + \Delta\theta} S_{1,2}(\theta) e^{-i x_{\text{size}} w (\theta_0 - \theta)} d\theta \quad (2.39)$$

where the indices 1 and 2 correspond to the TE and TM modes, respectively. The variable w denotes the projection of the sphere in the observation plane, and the complex amplitude of the scattered wave, $S_{1,2}(\theta)$, is derived from the LMT. The angle $\Delta\theta$ represents the opening angle of the imaging lens, which is determined by the aperture diameter and the distance between the lens and the sphere. Peaks in the intensity of $A_{\text{GP},1,2}(w)$ correspond to the glare points previously derived in the GO model.

2.3.2 Ray Transfer Matrix Model

The ray transfer matrix model (RTM), often referred to as the *ABCD model*, is a GO approach used to describe the propagation of light rays in paraxial optical systems [22, 60]. The RTM is commonly employed to summarise the behaviour of light in complex arrangements of optical elements, such as those encountered in DPTV and IPI systems. This section first introduces the general concept and definition of the RTM. It then describes RTMs for relevant optical elements, and concludes by explaining how complex optical systems can be represented using RTMs.

The RTM provides a simple framework for describing a 2D optical system. In this model, the z -axis defines the direction of the optical axis. A light ray at any point z_{ray} along the optical axis is characterised by two variables: its distance from the optical axis, x_{ray} , and its angle ϕ_{ray} relative to the optical axis. A light ray at position $z_{\text{ray},0}$ is therefore defined by the pair $(x_{\text{ray}}, \phi_{\text{ray}})_0$.




Figure 2.13: Notation for the ray transfer matrix system visualised. A ray is characterised by its distance and its angle to the optical axis

To describe the transfer of a ray $(x_{ray}, \phi_{ray})_0$ from position 0 to position 1 along the optical axis, a matrix M_{tot} is introduced. This matrix represents the total (*tot*) optical system between the two positions:

$$\begin{pmatrix} x_{ray,1} \\ \phi_{ray,1} \end{pmatrix} = \underbrace{\begin{bmatrix} A_{tot} & B_{tot} \\ C_{tot} & D_{tot} \end{bmatrix}}_{:= M_{tot}} \begin{pmatrix} x_{ray,0} \\ \phi_{ray,0} \end{pmatrix} \quad (2.40)$$

In Eq. (2.40), the matrix M_{tot} consists of four components: A_{tot} , B_{tot} , C_{tot} , and D_{tot} , hence the alternative name *ABCD matrix*. It should be noted that some literature uses the inverse notation (ϕ_{ray}, x_{ray}) , which results in a different arrangement of the matrix elements.

Ray Transfer Matrices for Relevant Optical Elements

Paraxial optical systems typically contain elements such as lenses, interfaces, and mirrors, all of which can be described using the RTM. In this section, only matrices for optical elements relevant to the later chapters are introduced. An overview of matrices for different optical elements can be found in Gerrard and Brunch [60]. The simplest case is the undisturbed propagation of a light ray through free space. The matrix for this case is denoted M_p (for *propagation*) and is given by

$$M_p = \begin{bmatrix} 1 & z_{01} \\ 0 & 1 \end{bmatrix} \quad (2.41)$$

where $z_{01} = z_{ray,1} - z_{ray,0}$ is the distance between the two considered positions along the optical axis.

The next relevant optical element is refraction at an interface, where n_i is the refractive index of the initial medium and n_j is the refractive index of the medium beyond the interface. The radius of curvature of the interface is denoted by R_I , where $R_I > 0$ indicates a convex interface. The refraction at the interface is described by

$$M_{I,i \rightarrow j} = \begin{bmatrix} 1 & 0 \\ \frac{n_i - n_j}{R_I n_j} & \frac{n_i}{n_j} \end{bmatrix}. \quad (2.42)$$

Later in this work, a notation will be used to indicate which media an interface separates. For example, $M_{I,w \rightarrow g}$ denotes a transition from water to glass.

In the case of a flat interface, where $R_I \rightarrow \infty$, the term $(n_i - n_j)/(R_I n_j) \rightarrow 0$ approaches zero, and the matrix simplifies to

$$M_{I,i \rightarrow j} = \begin{bmatrix} 1 & 0 \\ 0 & \frac{n_i}{n_j} \end{bmatrix} \quad (2.43)$$

A thin lens is described by

$$M_{\text{lens}} = \begin{bmatrix} 1 & 0 \\ 1 & -\frac{1}{f_{\text{lens}}} \end{bmatrix} \quad (2.44)$$

where f_{lens} is the focal length of the lens. A thick lens can be modelled as a combination of two curved interfaces and a propagation segment in between. There are many additional matrices for elements such as curved mirrors and prisms; however, these are beyond the scope of this work.

The Ray Transfer Matrix for Full Optical Systems

With the matrices for individual elemental components, complete systems can be described as a sequence of these elements. The total RTM of a system is given by

$$M_{\text{tot}} = \prod_{\ell} M_{\ell} \quad (2.45)$$

where the order of RTMs in the multiplication follows the path that the light travels. For example, if light travels from point 0 to point 2 with a thin lens located at position 1, the corresponding matrix is given by $M_{\text{tot}} = M_{p,12} M_{\text{lens}} M_{p,01}$.

There is also a physical interpretation for the individual RTM entries. In a lens system, $A_{\text{tot}} = x_{\text{ray},2}/x_{\text{ray},0}|_{\phi_{\text{ray},0}=0}$ describes the height change of a ray travelling through the system and thus represents the magnification of the optical system. The second entry, $B_{\text{tot}} = x_{\text{ray},2}/\phi_{\text{ray},0}|_{x_{\text{ray},0}=0}$, becomes zero when $x_{\text{ray},2} = 0$. This condition corresponds to the focus condition, and therefore $B_{\text{tot}} \neq 0$ describes the amount of defocus in an optical system. Using Eq. (2.44), it can be seen that the effective focal length of a system is given by $-1/C_{\text{tot}}$. The entry D_{tot} can be interpreted as the angular magnification of the system.

Beyond ray tracing, the RTM also finds application in wave optics, where it can be used to describe the propagation of waves through complex paraxial systems [142]. This is based on the interpretation that a light ray is always perpendicular to the wavefront of a spherical wave. Since the RTM models the propagation of any ray through an optical system, it can equivalently be used to describe the evolution of the corresponding wavefront.

2.3.3 Propagation of Light waves

While the GO model provides a useful framework for describing the propagation of light through an optical system, it does not fully account for diffraction and interference phenomena. Therefore, this section outlines the fundamental principles required to describe the propagation, diffraction, and interference of spherical waves within optical systems. Subsequently, the Rayleigh-Sommerfeld diffraction theory is presented to describe the propagation of spherical waves between two planes, such as the particle plane and the observer plane, capturing diffraction and interference effects, all of which are relevant in IPI. The RTM system is then combined with the Rayleigh-Sommerfeld diffraction theory to model wave propagation through complex paraxial optical systems, such as

camera lenses. The content of this section establishes the theoretical basis for Chapter 4, which develops a forward model used to simulate particle images.

Helmholtz Equation

Before introducing a model for wave propagation between two planes, it is necessary to define the type of waves under consideration and the underlying mathematical model. A wave in space and time is generally described by the full wave equation. In many cases, particularly with simple geometries and boundary conditions, this equation can be solved using the method of separation of variables, yielding both temporal and spatial components of the solution [19]. However, in optical measurement techniques employing standard camera sensors, light is integrated over a finite time interval (i.e. the exposure time). This process is equivalent to temporal averaging, which renders the temporal component of the wave negligible. Consequently, only the spatial solution is of interest. When focusing solely on the spatial component, the Helmholtz equation can be straightforwardly derived from the separation of variables approach by only considering the spatial solution [19]. The homogeneous form of the Helmholtz equation is given by:

$$\nabla^2 U(\varrho) + \kappa^2 U(\varrho) = 0 \quad (2.46)$$

where U denotes the scalar field at the 3D position $\varrho = (x, y, z)^T$, $\kappa = 2\pi/\lambda$ is the wave number, and $\nabla^2 U = \text{div}(\text{grad}(U))$ represents the Laplacian operator. The homogeneous Helmholtz equation arises from the full wave equation under the assumption of time-harmonic (single-frequency) dependence, i.e. $U(\varrho, t) = \Re\{u(\varrho)e^{-i\omega t}\}$. It describes source-free propagation of monochromatic waves, and is widely used in scenarios involving free-space propagation, diffraction, and interference [19].

Wave Propagation between two Planes

To describe the propagation of a spherical wave from one plane to another, various models exist. A widely used approach is the Rayleigh-Sommerfeld diffraction theory, which is particularly valuable as it is valid in both the near and far field. It can describe the propagation, diffraction, and interference of monochromatic spherical waves in homogeneous media [19]. The theory assumes the field value G to be scalar and incorporates Huygens principle, wherein every point on a wavefront acts as a point source emitting spherical waves, and the resulting field is the superposition of these waves.

For a known field G_0 in a source plane located at $z = 0$, the amplitude of the field G_1 at an arbitrary position (x_1, y_1, z_1) in another plane at $z = z_1$ can be computed using:

$$G_1(x_1, y_1, z_1) = \frac{1}{i\lambda} \iint G_0(x, y, z=0) \frac{z}{|\varrho|^2} e^{i\kappa|\varrho|} dx dy \quad (2.47)$$

where ϱ denotes connecting vector between the points $(x, y, z=0)$ at the source plane and the point (x_1, y_1, z_1) in the image plane, with the distance between points in the two planes being:

$$|\varrho| = \sqrt{(x - x_1)^2 + (y - y_1)^2 + z_1^2} \quad (2.48)$$

The derivation of the Rayleigh-Sommerfeld theory is lengthy and complex, and is thus omitted here. Further details can be found in Born and Wolf [19]. As solving the full Rayleigh-Sommerfeld integral is computationally intensive, simplified versions have been developed.

One such simplification is Fresnel's approximation, which assumes that the separation between the planes is large compared to their lateral extent in (x, y) . Under this assumption, two simplifications can be made. When z is large relative to the in-plane distances $x - x_1$ and $y - y_1$, the approximation $\frac{1}{|\varrho|^2} \approx \frac{1}{z_1^2}$ can be applied to the prefactor in the integral. However, this simplification cannot be applied to the exponential term, as even small variations in phase can significantly affect the result. Instead, a Taylor expansion is used in Fresnel's approximation:

$$|\varrho| \approx z_1 \left(1 + \frac{(x - x_1)^2 + (y - y_1)^2}{2z_1^2} \right). \quad (2.49)$$

Substituting this approximation into the Rayleigh–Sommerfeld integral, the Huygens–Fresnel integral is obtained:

$$G_1(x_1, y_1, z_1) = \frac{e^{i\kappa z_1} e^{i\kappa \frac{x_1^2 + y_1^2}{2z_1}}}{i\lambda z_1} \iint G_0(x, y, z = 0) e^{i\kappa \frac{x^2 + y^2}{2z_1}} e^{-\frac{\kappa i}{z_1}(x_1 x + y_1 y)} dx dy. \quad (2.50)$$

Generalization to Paraxial Optical Systems

The Huygens–Fresnel integral can be used to model wave propagation between two planes through free space. However, when modelling wave propagation through optical systems relevant to IPI and DPTV, optical elements such as lenses, interfaces, and aperture stops must be considered. This is where the RTM, introduced in Section 2.3.2, becomes relevant. By incorporating the RTM into the Huygens–Fresnel integral under the paraxial approximation, a generalised formulation can be derived [15, 36]:

$$G_1(x_1, y_1, z = z_{\text{eff}}) = -i \frac{\kappa}{2\pi \sqrt{B_{\text{tot}}^x B_{\text{tot}}^y}} \exp(i\kappa z_{\text{eff}}) \iint G_0(x, y, z = 0) \exp\left(-\frac{i\kappa}{2B_{\text{tot}}^x} [D_{\text{tot}}^x x^2 - 2x_1 x + A_{\text{tot}}^x x_1^2]\right) \exp\left(-\frac{i\kappa}{2B_{\text{tot}}^y} [D_{\text{tot}}^y y^2 - 2y_1 y + A_{\text{tot}}^y y_1^2]\right) dx dy. \quad (2.51)$$

Here, two distinct RTMs are considered: one for the (x, z) plane $(A_{\text{tot}}^x, B_{\text{tot}}^x, D_{\text{tot}}^x)$ and one for the (y, z) plane $(A_{\text{tot}}^y, B_{\text{tot}}^y, D_{\text{tot}}^y)$. This approach is particularly useful when dealing with optical elements such as cylindrical lenses or interfaces that are curved in one plane but not the other (e.g., a cylindrical wall). The variable $z_{\text{eff}} = \sum_{\ell} z_{\ell}/n_{\ell}$ denotes the effective optical distance between the two planes, accounting for refractive indices n_{ℓ} along the way.

Eq. (2.51) enables the modelling of light wave propagation through paraxial optical systems, including lenses and other optical components. However, this formulation does not account for aperture stops. A method to incorporate aperture stops was proposed by Yura and Hanson [221], in which the wave is propagated from the initial plane to the aperture stop plane, where the field is truncated. The field in subsequent planes is then computed from this truncated field.

2.4 Relevant Basics in Deep Learning for Particle Detection, Classification and Image Generation

This section briefly introduces fundamentals of deep learning that form the theoretical foundation for the methods used in Chapter 5 to detect PIs and generate images for training. First, Section 2.4.1 presents the concept of neural networks. Then, Section 2.4.2 outlines general concepts in the training of neural networks and deep learning. Following that, Section 2.4.3 discusses convolutional neural networks, the central type of neural networks relevant to this thesis. Subsequently, Section 2.4.4 introduces selected neural networks for object detection, which are relevant to the detection of PIs in optical measurements. Finally, Section 2.4.5 presents a type of neural network relevant to image generation, as used in Chapter 5.

In deep learning, a transfer function is fitted to data through a training process to solve an abstract task, without the need to explicitly formulate the transfer function or algorithm as in fixed programmes. This approach enables machine learning algorithms to address complex tasks that are too difficult for fixed programmes [66]. Such algorithms are widely applied in computer vision [164, 120, 104], image generation [67, 96, 76], speech recognition [77, 69], and natural language processing [147, 196]. With the rise of machine learning driven by advancements in graphics processing units, models have proliferated across various fields. Nevertheless, the field of deep learning remains to be theoretically fully understood due to the high-dimensional and non-convex nature of the optimisation process.

2.4.1 Neural Networks

Neural networks consist of a multitude of neurons that are interconnected in a graph. These neurons are typically arranged in layers, which are stacked on top of each other. The resulting graph represents a function capable of mapping an input to a more abstract output. The parameters of this function (i.e. the graph) are updated during training in order to accomplish a specific task. A typical structure of a neural network is illustrated in Fig. 2.14.




Figure 2.14: Schematic structure of a fully connected feed forward neural network. There is an input layer x_i , two hidden layers h_i and an output layer y_i . Figure adapted from [172]

A neural network generally begins with an input layer of neurons into which the input is directly fed. The size of the input layer depends on the nature of the input. For instance, a single number as input would require only one input neuron, whereas an image of size $i \times j$ with three colour channels would necessitate an input layer of size $i \times j \times 3$.

The input layer then passes the information to the hidden layers. These hidden layers enable the network to fit complex non-linear mappings to the output. The size and configuration of the hidden layers can vary, and the task of identifying suitable architectures for such network graphs remains an active area of research. The hidden layers are responsible for transforming the input into the desired output within the neural network.

Following the hidden layers, the output layer is located at the end of the network. Similar to the input layer, the shape of the output layer corresponds to the specific task the network is designed to perform. For example, a regression task aiming to predict a single integer would require only one output neuron, whereas a classification task would need ℓ neurons for ℓ classes. If the task involves generating an image, the output layer would be shaped accordingly to match the image dimensions.

Neurons

The building blocks of a neural network are the individual neurons. The structure of a neuron is illustrated in Fig. 2.15. A neuron comprises a linear function followed by a non-linear activation function. The linear component of the neuron assigns a weight ω_i to each input x_i , sums the weighted inputs, and adds a bias term b :

$$z = b + \sum_i x_i \omega_i \quad (2.52)$$

In a hidden layer, the inputs x_i are the outputs of neurons from the previous layer. The weight ω_i thus represents the strength of the connection between a neuron in the current hidden layer and a neuron in the preceding layer. It can be interpreted as the degree to which a neuron responds to the inputs from other neurons.




Figure 2.15: Schematic structure of a neuron in a neural network. The neuron receives inputs x_i which are weighted by the trainable weights ω_i . Each input is weighted and then summed up and a bias b is added. The results is then fed through an activation function $a(z)$ to produce the final output of the neuron y . Figure adapted from [172].

The weights ω_i and the bias b constitute the learnable parameters of a neural network and are updated during the training process. These parameters are initially set at random and iteratively refined throughout training.

In a fully connected network, each neuron in the current layer is connected to all neurons in the previous layer. A fully connected network comprising an input layer, multiple hidden layers, and an output layer is referred to as a multilayer perceptron (MLP). According to the universal approximation theorem [38, 78], a foundational result in neural network theory, a feed-forward

neural network with at least one hidden layer can approximate any continuous function to an arbitrary degree of accuracy, provided it contains a sufficient number of neurons in the hidden layer.

Activation Function

In order for a neural network to approximate complex functions, non-linearity must be introduced. A combination of neurons composed solely of linear functions would result in an affine transformation. The activation function $a(z)$ introduces non-linearity, thereby enabling the model's capacity to increase with the number of neurons and allowing it to fit more complex distributions.

The four most commonly used activation functions are the hyperbolic tangent (Eq. (2.53)), the sigmoid function (Eq. (2.54)), the rectified linear unit (ReLU) (Eq. (2.55)), and leaky ReLU (Eq. (2.56)). The hyperbolic tangent is typically used in hidden layers, as its output range of $[-1, 1]$ helps to centre the output values. The sigmoid function is commonly used in the output layer for binary classification problems, as it produces values in the range $[0, 1]$, which can be used as class probabilities.

$$\tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}} \quad (2.53)$$

$$\sigma(z) = \frac{1}{1 + e^{-z}} \quad (2.54)$$

$$\text{ReLU}(z) = \begin{cases} 0 & \text{if } z \leq 0 \\ z & \text{if } z > 0 \end{cases} \quad (2.55)$$

$$\text{leaky ReLU}(z) = \begin{cases} cz & \text{if } z \leq 0 \\ z & \text{if } z > 0 \end{cases} \quad (2.56)$$

The ReLU activation function is frequently used in hidden layers and offers several advantages over exponential activation functions. Firstly, ReLU enhances computational efficiency, as computing $\max(z, 0)$ is less intensive than evaluating exponential functions [140]. Secondly, ReLU promotes sparsity in the network due to its zero output for negative inputs. Lastly, ReLU accelerates training convergence, as it does not saturate for large values of z , unlike exponential activation functions [140]. A major drawback of ReLU is that it outputs zero for any negative z value. This leads to a zero gradient during back-propagation, preventing the weights of the neuron from being updated. This issue, known as the dying ReLU problem, is addressed by leaky ReLU, which introduces a small negative slope c for negative z values [125].

2.4.2 Training of Neural Networks

There are three paradigms of machine learning: supervised learning, unsupervised learning, and reinforcement learning. The first learning paradigm is supervised learning, in which the model learns from labelled data, where the network's prediction is compared to the ground truth (labels). Supervised learning is often used for regression and classification models, such as in object detection. It is a straightforward approach to implement, but typically requires large amounts of labelled data. In unsupervised learning, the second learning paradigm, the network learns patterns and structures from unlabelled data without specific guidance. This method is often used in image generation models. Unsupervised learning bears the advantage of not requiring pre-labelled data, but can be more complex to realise. The third learning paradigm is reinforcement learning, in which the network learns from interactions with the environment, receiving rewards or penalties based on its actions. Reinforcement learning is often used to train agents and decision-making models [190].

Despite the existence of different approaches to machine learning, the common basis of all learning paradigms is the gradient-based update of the weights. The process typically consists of three steps:

the forward run, where the network outputs a prediction; the computation of the loss function; and the update of the weights by back-propagation.

Learning Iteration: Forward Run, Loss Computation and Backpropagation

In the first step, the forward run, the network makes a prediction based on the input and the network's current weights in the network graph. In supervised learning, the prediction y is compared to the ground truth label \hat{y} , and the deviation between the two is computed.

This deviation is quantified by a loss function $L(y, \hat{y})$, which typically serves as the objective function to be optimised during training. The goal of training is to minimise the loss function. Typical loss functions include mean square error (Eq. (2.57)) and L2 loss, used for regression tasks, and the cross-entropy loss (Eq. (2.58)), used for classification tasks.

$$L_{\text{MSE}}(y, \hat{y}) = \frac{1}{N} \sum_{\ell=1}^N (y_{\ell} - \hat{y}_{\ell})^2 \quad (2.57)$$

$$L_{\text{CE}}(y, \hat{y}) = -\frac{1}{N} \sum_{\ell=1}^N [y_{\ell} \log(\hat{y}_{\ell}) + (1 - y_{\ell}) \log(1 - \hat{y}_{\ell})] \quad (2.58)$$

With the loss function computed, the weights are updated in a process called backpropagation [168]. In backpropagation, the gradient of the loss function with respect to each weight in the network is computed using the chain rule. The weights are then updated based on the gradient, multiplied by a factor called the learning rate. The gradient provides the direction of change for the weights, and the learning rate determines the magnitude of the update. A low learning rate can cause unnecessarily slow convergence of the training process and may result in the optimiser getting stuck in a local minimum. A high learning rate, on the other hand, can cause the network to overshoot a minimum and lead to oscillations. The learning rate is one of the most crucial hyperparameters to choose for training (hyperparameters are parameters that are not updated during training).

The three steps, forward run, loss computation, and backpropagation, are then iterated until the process converges. Since the gradient is computed with respect to the neurons weights, the activation functions in the network must be continuously differentiable. Further details on the training of neural networks can be found in [66] and [110].

Optimisers for Training

In order to find a minimum of the loss function, a high-dimensional non-convex optimisation problem must be solved. To obtain a suitable solution in this high-dimensional weight space, different optimisers are employed. The most straightforward approach is full gradient descent, in which the model is applied to all data in the dataset and a gradient is computed for the entire dataset. However, this approach quickly becomes unfeasible for large datasets. Instead, only a mini-batch of the full dataset is selected at random, and the gradient for the optimisation is computed based on the result for this mini-batch. This approach is called stochastic gradient descent (SGD). SGD has the advantage of requiring significantly less computation for a weight update compared to full gradient descent. However, the method assumes that the randomly sampled mini-batch is representative of the entire dataset. This means that the mini-batch size must be adjusted based on the diversity of the full dataset to ensure a high probability that the mini-batch sufficiently represents the whole dataset. The dataset is split into a number of equally sized mini-batches, and

at each optimisation step, one of these mini-batches is selected at random. Once all mini-batches have been processed, one epoch is completed, and the process begins again with the training dataset being split into new mini-batches. The mini-batch size is, therefore, another important hyperparameter to select during training.

An improved method of SGD is SGD with momentum (SGDM) [152], which aims to help the optimiser avoid getting stuck in local minima. In this approach, a fraction of the gradient from the previous iteration is added to the current weight update to ‘keep the momentum’ of the optimisation. The approach is inspired by the way a sphere would roll down a hill under the influence of gravity and momentum conservation. The momentum can help to escape local minima and dampen rapidly changing gradients. The strength of the momentum in SGDM relative to the current gradient is also a hyperparameter.

One of the most widely adopted extensions of SGDM is the adaptive momentum estimation (ADAM) [95]. This approach combines the advantages of both momentum and adaptive learning rates by maintaining exponentially decaying averages of past gradients (first moment) and squared gradients (second moment). These estimates are then bias-corrected to account for their initialization at zero. ADAM adapts the learning rate for each parameter individually, making it particularly effective for problems with sparse gradients or noisy data.

Further Considerations on Training

There are two types of error in training neural networks that need to be considered: the optimisation error and the generalisation error. The optimisation error is the error that is directly reduced during training. It reflects how well the network predicts the training data. As this is a high-dimensional, non-convex problem, it is unlikely that the global minimum will be found during training. Instead, the process typically converges to a local minimum, or becomes stuck at a saddle point or plateau - i.e. an optimisation error remains. The optimisation error only captures the network’s performance on the training dataset. The generalisation error, on the other hand, measures the error of the neural network when applied to a new, previously unseen dataset. The generalisation error is the relevant metric when the network is later used in inference. However, it cannot be directly minimised. The generalisation capability of a network usually depends on how well the training data represents real data and on the capacity of the model. The process by which a model begins to memorise the training data instead of learning the underlying features is called overfitting [66]. Overfitting occurs based on the relationship between the network’s capacity and the diversity and size of the training set. Model capacity refers to the model’s ability to fit complex distributions and generally increases with the size of the network. The size and diversity of the training data determine how easily a network of a given capacity can overfit the dataset. When overfitting occurs, the generalisation error typically increases, impairing the network’s performance in real-world applications.

However, the opposite case can also occur when the model lacks the capacity to sufficiently fit the training dataset. This is known as underfitting [66]. Therefore, a proper balance between model capacity and training data must be achieved.

While the optimisation loss is directly measured by the loss function, the generalisation error can be estimated using the validation loss. The validation loss is obtained as follows: the total dataset is split into two parts, the training dataset (typically 80-90% of the total set), which is used during training, and the validation set (typically 10-20% of the total set). The validation set is withheld from the network during training. At regular intervals, the network is applied to the validation set, and the same loss function used for the optimisation loss is used to compute the validation loss.

The key difference is that the validation step does not influence backpropagation, as the network should not be trained on the validation set. The validation loss thus provides an estimate of how well the neural network generalises to new, unseen data. Since the computation of the validation loss does not contribute to training, the validation interval should be chosen to be sufficiently large to avoid unnecessary computation. The validation interval is another hyperparameter.

During training, both the optimisation loss and the generalisation loss should be minimised. This is usually achieved through regularisation and by monitoring the validation loss. At the beginning of training, both losses typically decrease as the network begins to learn the data representation. However, at some point, the network may begin to overfit, and the validation loss will plateau or even increase. A divergence between the training and validation loss is a typical sign of overfitting. A method to prevent this is early stopping, where training is terminated as soon as the validation loss increases. However, stopping training immediately when the validation loss increases may be premature, as the validation loss can fluctuate due to noise or variance in the data. To address this, a technique known as early stopping with patience is employed [148]. In this approach, the model's parameters are saved whenever a new lowest validation loss is achieved. Training continues, and only terminates if the validation loss does not improve for a specified number of consecutive evaluation steps (the patience parameter). At the end of training, the weights corresponding to the lowest recorded validation loss are restored.

Other methods to avoid overfitting include dropout regularisation [187] and L2 regularisation. In dropout, a certain number of neurons are disregarded during the current iteration. This forces the network to form new interconnections between different neurons. L2 regularisation sums the weights of all neurons and penalises them for becoming too large. Other commonly used techniques include batch normalisation [81] and learning rate decay.

2.4.3 Convolutional Neural Networks

As this work is centred around optical measurement techniques and the associated image processing, this section focuses on neural networks specifically designed for image-related tasks.

Convolutional Neural Networks (CNN) [108, 112] are a widely adopted neural network architecture for processing images, videos, and other data with a known grid-like structure [66]. Since the success of AlexNet [103, 104] in the 2012 ImageNet competition [44, 169], deep CNNs have become the state-of-the-art architecture in computer vision. With the advent of the transformer architecture [203], vision transformers [47] have emerged as an alternative in the field of computer vision. However, CNNs remain widely used, as transformers require significantly larger training datasets and computational resources, and offer only limited advantages when contextual information is not essential.

Convolutional Layers

The core concept of CNNs is the convolutional layer, which has the significant advantage of weight sharing. This allows it to process images using far fewer weights than fully connected layers would require. Convolutional layers perform spatial convolution as defined in Eq. (2.59).

$$y[u, v] = \sum_{\ell=1}^{\infty} \sum_{k=1}^{\infty} h[\ell, k] \cdot x[u - \ell, v - k] \quad (2.59)$$

A convolutional layer uses a kernel $h[\ell, k]$, also called a filter, which slides over the input x with a predetermined stride. At each stride step, the kernel computes the convolution value based on the current input region to produce the output $y[u, v]$ of the layer. This operation is repeated for every stride step, as illustrated in Fig. 2.16.

This process can be interpreted as pattern recognition: if the filter and the input patch are similar, the filter will produce a high output value. Analogous to Fig. 2.15, a bias can be added to the filter output, followed by the application of a non-linear activation function. When the CNN learns the parameters of a kernel, the kernel effectively becomes a filter for a specific feature that the network has learned.




Figure 2.16: Spatial convolution visualised: A 3×3 kernel (gray cells) is moved over a 5×5 input (green cells) with padding (white cells). The kernel moves with a stride of one producing a 5×5 output (blue cells). Figure adapted from [49].

A convolutional layer has the advantage that its parameters are confined to the kernel, and these parameters are reused across the entire input. A convolutional layer can contain multiple kernels, each corresponding to a different feature, such as edges, shapes, or textures. Each kernel produces a separate output layer, known as a feature map. These maps indicate how strongly a particular feature (i.e., the activation of the corresponding kernel) is present at each position in the input. The number of kernels in a convolutional layer, hence the number of feature maps it produces, is defined by a hyperparameter of the network architecture.

Pooling Layers

The other fundamental building block in CNNs is the spatial pooling layer. Its primary purpose is to reduce the size of the output feature map and decrease the number of parameters in the CNN, thereby lowering computational requirements and potentially reducing the risk of overfitting. Spatial pooling operates in a manner similar to convolutional layers, using a fixed-size kernel that slides across the input based on a defined stride. However, instead of performing a convolution operation, pooling layers apply different functions, such as taking the maximum value within the considered region (max pooling) or computing the average value (average pooling).

"Max pooling outputs the maximum value in the considered neighbourhood, highlighting the most prominent features of the input at the cost of losing information on less strong ones. Average pooling computes the average value of the considered input patch, smoothing the input but at the cost of losing information on strong features." [172]. Among the various pooling operations, max pooling has emerged as the most commonly used, as it provides the best performance in practice [175]. In contrast to convolutional layers, the parameters of pooling layers are fixed and non-learnable.

Architecture of a simple CNN

A simple CNN consists of multiple sequential blocks of convolutional and pooling layers. In many CNNs, the spatial dimensions of the convolutional layers decrease in height and width, while their depth increases as the information propagates deeper into the network, compare the boxes representing feature maps in Fig. 2.17. As the depth increases, the features become more semantically rich and abstract. Understanding exactly what features CNNs learn remains an open research question. Various approaches in explainable AI, such as Score-CAM [207], aim to visualise these learned features. This part of the CNN can be regarded as the feature extraction stage.

At the end of the sequential convolutional and pooling layers, the resulting feature maps are typically flattened into a single long array that contains all the abstract features extracted by the convolutional layers. This feature vector is then passed into a fully connected layer, which acts as the transfer function for the desired task, such as classification or regression, similar to a MLP. The structure of a simple CNN is illustrated in Fig. 2.17.




Figure 2.17: Simple architecture of a CNN classifier. An input image is processed using a convolutional layer with multiple filters. The network used multiple convolutional layers, decreasing in size but increasing in depth, increasing the level of abstraction of features. The final feature vector is flattened and processed by a two fully connected layers to perform the classification tasks. Figure adapted from [130]

Residual Neural Networks

The performance of CNNs can be enhanced by incorporating more complex architectural structures. Such structures help to address challenges like vanishing and exploding gradients during training, which are central issues in deeper networks. Deeper CNNs have been shown to yield improved accuracy on image processing tasks [184, 191].

One widely adopted method is the inclusion of residual blocks, which contain skip connections [72], as illustrated in Fig. 2.18. These skip connections introduce cross-layer dependencies and help mitigate the vanishing gradient problem by allowing gradients to flow more directly through the network. This enables CNNs to be significantly deeper without suffering from training degradation. Popular architectures that employ this approach include ResNet50 and ResNet101 [72].

2.4.4 Object Detection Networks

CNNs are the basis for many object detection networks, which are specialised networks typically used to perform the detection and classification of multiple objects in an image. These kinds




Figure 2.18: Residual block: The output from a previous layer is added to the output of the current layer. Figure adapted from [72]

of networks are especially helpful for the tasks of particle detection in methods such as DPTV [SDLK23, 48, 99, 35]. Object detection networks often have specialised architectures that allow them to perform the detection and classification tasks simultaneously. However, most of them have in common that they employ a CNN-based feature extraction backbone (e.g. ResNet50) at the beginning of the network, to feed the rest of the network.

In object detection, there are typically two kinds of tasks: localisation and classification:

- **Localization (where question):** The localisation task typically involves placing a bounding box (BB) over an object to characterise the object's position and size (width and height). The localisation task can be formulated as a regression task of four parameters, e.g. BB centre position (x_{BB}, y_{BB}) and the BB height and width w_{BB}, h_{BB} , i.e. $BB = (x_{BB}, y_{BB}, w_{BB}, h_{BB})^T \in \mathbb{R}^4$.
- **Classification (what question):** The classification task involves predicting the semantic class of detected objects, i.e. a discrete categorisation problem. This is typically done by outputting a probability vector $p_C = (p_1, p_2, \dots, p_N) \in \mathbb{R}^{N_C}$, with $|p_C| = 1$, which contains a probability for all N_C classes, for the present object to be in the class C_ℓ , $\ell = 1, \dots, N_C$. Usually, the outputted class is the entry in p_C with the highest probability.

Object detection networks are typically distinguished into two architecture-specific implementations:

- **Two-stage detectors (e.g. Faster R-CNN [164]):** Two-step detectors usually first generate region proposals for an object. This involves a binary classification (object/background) and a coarse regression of the region of interest (RoI). Only in the second step is the object in the RoI then localised (refined regression) and classified (multi-class classification).
- **Single-stage detectors (e.g. YOLO [163]):** A simultaneous multi-class classification and BB regression is performed in one forward pass using a unified framework.

Single-stage detectors are usually used for tasks such as the real-time detection of objects, where the speed during inference is of the essence. Multi-stage detectors are usually optimised for accuracy rather than speed.

In the following, two popular object detectors, Faster R-CNN [164] and YOLO [163], are introduced, as these two architectures are used in Chapter 5.

Region Based Convolutional Neural Networks

One of the most used two-stage detectors is Faster R-CNN [164]. It is the successor of Fast R-CNN [62], which in turn was developed from Region-based CNN (R-CNN) [63].

R-CNN introduced the concept of combining region proposals with CNNs for object detection. The approach operates in three stages: First, a conventional algorithm such as selective search [199] generates approximately 2,000 region proposals per image. Each region proposal is then individually processed by a CNN to extract features and classify the object. Finally, features are fed into a fully connected layer to perform the BB regression task.

To enable training of an object detection network like R-CNN, labelled ground truth BBs are required, specifying the actual position and class of each object in the image. In the R-CNN framework, the regressor learns to map region proposals to these ground truth BBs. During training, each predicted BB is matched with the closest ground truth BB based on their intersection over union (IoU), which quantifies the overlap between the two. More detail on the IoU is provided in Chapter 6.

R-CNN suffers from computational inefficiency, as each region proposal is fed through the CNN separately, resulting in redundant feature computation. Fast R-CNN addresses this issue by sharing the convolutional computation in the CNN. The entire input image is passed through the CNN to generate a feature map. Selective search is applied on the original input image, and the region proposals are projected onto the feature map outputted by the CNN. These RoIs usually have variable size and shape. As the following fully connected layers require fixed-size inputs, RoI pooling is applied to convert the variable-sized RoIs into a standard size that can be processed by the rest of the network. The sequence of fully connected layers bifurcates into two output layers: one for BB regression and one for classification.

Fast R-CNN employs a multi-task loss combining a classification loss L_C , typically cross-entropy, with a BB regression loss L_{BB} , which is a smooth L1 loss [62].

$$L_{\text{multitask}} = L_C + cL_{BB} \quad (2.60)$$

The importance of the class loss versus the BB loss is regulated using a hyperparameter c . The smooth L1 loss smooths out the sharp edge of L1 at $|y - \hat{y}| = 0$.

$$\text{smooth L1}(y, \hat{y}) = \begin{cases} 0.5|y - \hat{y}|^2 & \text{if } |y - \hat{y}| < 1 \\ |y - \hat{y}| - 0.5 & \text{otherwise} \end{cases} \quad (2.61)$$

While Fast R-CNN improved the computational efficiency of R-CNN, it still relies on external algorithms such as selective search for region proposals. Faster R-CNN addressed this bottleneck by introducing a Region Proposal Network (RPN), achieving a fully end-to-end trainable object detection system. In Faster R-CNN, a large CNN backbone (e.g. ResNet50) processes the full image to extract a feature map, see Fig. 2.19. The RPN then takes this feature map as input. It uses anchor boxes, which are uniformly distributed over the feature map. The idea behind anchor boxes is that the network does not need to predict BBs from scratch, but instead refines standard-sized anchor boxes of various sizes and aspect ratios. The RPN slides over the image and processes these anchor boxes. For each anchor, a BB regression is performed along with a classification (object/background). Non-maximum suppression is applied to remove overlapping proposals and avoid duplicate detections.




Figure 2.19: Structure of Faster R-CNN: The backbone, the RPN and the detection head are marked, as well as their tasks are pointed out. Modified figure from [30].

The proposals with the highest objectness scores (i.e. the probability of an object being present) are then passed to the second stage, the detection head. The detection head uses the output feature maps from the backbone and focuses only on the proposals. The detection head again uses a series of fully connected layers, which bifurcate into a multi-class classification layer and a BB regression layer. It determines the class of the RPN proposals and further refines the BBs.

You Only Look Once

One of the most popular examples of a single-stage detector is You Only Look Once (YOLO) [163]. YOLOv1 represents a different approach compared to Faster R-CNN. Unlike two-stage methods that rely on region proposals, YOLOv1 processes the entire image in one forward pass of a CNN. In this forward pass, BBs and class probabilities are predicted simultaneously. This unified architecture allows for shorter inference times, making the network particularly interesting for real-time applications.

The algorithm divides the input image into an $N \times N$ grid. Each grid cell is responsible for detecting objects within its boundaries. For each grid cell, a fixed number of BBs is predicted, each with a confidence score. The confidence score reflects the probability that the BB contains an object and how well the BB fits the object, as measured by the IoU. Additionally, class probabilities are predicted for each grid cell. For each BB, the final class-specific confidence score is computed as the product of the confidence score and the class probability. This score reflects both the likelihood that a BB contains an object and the likelihood of the object's class. BBs with low confidence scores are discarded, and non-maximum suppression is used to eliminate overlapping BBs referring to the same object.

During training, YOLOv1 learns to predict BB coordinates, objectness scores, and class probabilities by minimising a combined loss. This loss includes a localisation loss, which penalises BB placement errors, a confidence loss for incorrect objectness scores, and a classification loss. YOLOv1 employs 24 convolutional layers followed by two fully connected layers. The convolutional layers serve the purpose of feature extraction, and the fully connected layers perform the final regression of BB coordinates and class probabilities. By processing the full image at once,

YOLO helps reduce the number of false positives in the background and allows the network to learn more generalisable features, as objects are seen in context. However, as each grid cell predicts only one set of class probabilities, YOLOv1 struggles to detect small or closely spaced objects that fall within the same cell. Additionally, the coarse grid-based localisation can lead to less precise BB predictions compared to region-based approaches.

YOLOv2 (also known as YOLO9000) [161] introduced the concept of anchor boxes to YOLO. Unlike YOLOv1, where BBs were predicted from arbitrary shapes, which could lead to instability in training due to large differences to the ground truth, YOLOv2 predicts offsets from the predefined anchor boxes. In training, a preprocessing step groups together the ground truth BBs into similar shapes by k-means clustering, and the five most common shapes are selected as anchor boxes. The predicted offset of BBs from the anchor boxes is constrained to stabilise training and guide adaptation to object shapes. To support this change, the two fully connected layers were replaced with convolutional layers that directly output BB coordinates and class probabilities. Additionally, a pass-through layer was introduced that concatenates higher-resolution feature maps from earlier layers with deeper semantic features from later layers. This helps the network to detect small objects more efficiently by preserving higher spatial resolution.

YOLOv3 [162] improved the detection of differently sized objects. In a CNN, the spatial resolution usually gets reduced with every layer, which can cause problems resolving small objects. Inspired by the concept of feature pyramid networks (FPN) [119], YOLOv3 performs detections on three different feature maps. The last and coarsest feature map is used to detect large objects. This feature map is then upsampled and added to a feature map from earlier layers, allowing detection on a finer scale. This process is repeated to produce three detection scales, enabling the network to detect small, medium, and large objects using feature maps of different resolutions. Objects are detected on each of these feature maps, and nine different anchor boxes, three per feature map, are used. Additionally to these changes, YOLOv3 introduces a larger backbone, Darknet-53 [162], compared to the Darknet-19 used by YOLOv2 [161]. Darknet-53 is a CNN consisting of 53 convolutional layers with alternating 3×3 and 1×1 convolutions, as well as residual connections, similar to ResNet.

YOLOv4 [17] introduced two major improvements on YOLOv3. First, a new backbone CSP-Darknet-53 [17] was introduced, which is based on Darknet-53 and adds cross stage partial (CSP) connections [206]. CSP splits a feature into two parts, one of which processes further through the layers while the other bypasses these layers. The two feature parts are then merged before being processed further. By not processing the entire feature map, computational cost is reduced, which allows for deeper networks without a proportional increase in computational cost. The bypassing of information helps to improve the gradient flow, which can help to prevent exploding or vanishing gradients. The merging of processed and unprocessed feature maps forces the network to learn complementary feature representations, improving generalisation capability.

Another architectural change in YOLOv4 is the inclusion of a path aggregation network (PANet) [122] in its architecture. Unlike traditional FPNs, which only pass information top-down, PANet adds a bottom-up path. Top-down-only FPNs may not effectively propagate low-level spatial details, which are crucial for detecting small objects. PANet uses lateral connections to merge features from different levels in the CNN. The lower spatial resolution but more semantically rich feature maps are upsampled and merged with the high-resolution ones, to obtain high-resolution and semantically rich feature maps.

The last important change in the YOLOv4 architecture is the inclusion of spatial pyramid pooling (SPP) [71]. SPP is applied to the final feature map, after the feature extraction backbone and PANet, to capture features at multiple spatial scales and improve the detection of objects of various

sizes. It performs pooling operations with various kernel sizes (e.g. 1×1 , 5×5 , 9×9 , and 13×13) on the final feature map, which allows the network to gather both fine-grained (1×1) and contextual information (13×13) from the feature map. The pooled features are then concatenated in the channel dimension to increase the receptive field.

Further versions of YOLO, specialised for specific tasks, have been developed [83], however, YOLOv4 remains a widely used detector for a broad range of applications.

2.4.5 Generative Adversarial Networks

Neural networks can be used not only to detect objects in images but also to generate images. A widely used framework for image generation is the generative adversarial network (GAN) [67]. One of the key advantages of GANs is that they can be trained in an unsupervised manner, making them suitable for scenarios with limited labelled data.

A GAN consists of two networks, a generator and a discriminator, trained in an adversarial setup. The generator learns to produce synthetic images that resemble those in the training dataset, while the discriminator attempts to distinguish between real images and those generated by the generator. During training, the generator's objective is to fool the discriminator by generating images that are indistinguishable from real ones. Through this adversarial process, both networks improve over time. GANs are used in Chapter 5 to generate images to train object detectors.

Generator

The generator network learns a function that maps a random noise vector, sampled from a latent space — a compressed, abstract representation of data where hidden features and patterns are encoded — to an image that resembles images from the training data. This noise vector encodes abstract features such as shapes and textures. The generator upsamples the noise vector into a full-sized image using a series of transposed convolutional layers. Variations in the input noise vector lead to a variety of generated output images.

Transposed convolutional layers perform the inverse operation of standard convolutional layers. While a convolutional layer slides a kernel over the input to sum up the weighted values, a transposed convolution spreads out each input over a larger area and adds overlapping contributions together. Like standard convolutions, transposed convolutions use learnable kernels that are updated during training.

Discriminator

The discriminator is a CNN that functions as a binary classifier, predicting whether an input image is real or generated. During training, the discriminator learns features that characterise real images. As the generator's images become more realistic, the discriminator must learn increasingly refined features to distinguish fake samples from real ones, and this adversarial process continues iteratively.

Training of Generative Adversarial Networks

During training, both the generator and the discriminator are initialised with random weights. The training process consists of two alternating steps: training the discriminator and training the generator.

Discriminator Training: In the first step, the generator produces a batch of fake images, $\mathcal{G}(x_{\text{noise}})$, from random noise x_{noise} , while another set of real images, x_{real} , is sampled from the training dataset. These two sets are combined, labelled accordingly (real or fake), and passed to the discriminator. The discriminator then predicts the probability that each image is real. The discriminator's outputs on real data, $\mathcal{D}(x_{\text{real}})$, and on generated data, $\mathcal{D}(\mathcal{G}(x_{\text{noise}}))$, are used to compute the discriminator loss, following the cross-entropy loss:

$$L_D = -[\log(\mathcal{D}(x_{\text{real}})) + \log(1 - \mathcal{D}(\mathcal{G}(x_{\text{noise}})))]. \quad (2.62)$$

This loss is then used to update the discriminator's parameters via backpropagation.

Generator Training: In the second step, the generator is trained. A new batch of images, $\mathcal{G}(x_{\text{noise}})$, is generated from fresh noise vectors. These images are passed to the discriminator, which predicts whether the images are real or fake, $\mathcal{D}(\mathcal{G}(x_{\text{noise}}))$. The generator loss [67]

$$L_G = -\log(\mathcal{D}(\mathcal{G}(x_{\text{noise}}))) \quad (2.63)$$

is then computed as a measure of the generator's ability to fool the discriminator. The generator's parameters are subsequently updated using backpropagation of the gradient through the discriminator, while the discriminator's weights remain fixed.

These two steps are repeated over many iterations, ideally converging to a Nash equilibrium [141], where the discriminator can no longer distinguish fake from real images [67]. At this point, the discriminator outputs a probability close to 0.5 for both real and generated images, indicating that it is effectively guessing.

Training GANs is formulated as a min-max optimisation problem, as originally proposed by Goodfellow *et al.* [67]. This optimisation framework, however, can result in unstable training dynamics. One common issue arises when the discriminator overpowers the generator and easily identifies fake images. In this case, the generator receives vanishing gradients, meaning the feedback it obtains is too weak to guide learning effectively. As a result, the generator struggles to improve [6]. Another common issue is mode collapse, where the generator learns to produce only a limited set of outputs that consistently fool the discriminator. While these outputs may appear realistic, they lack diversity and fail to represent the full distribution of the training data [186]. This limits the generator's ability to generalise and reduces the quality of the generated data.

Building upon the foundational GAN framework, numerous architectural advancements have been proposed to improve training stability and image quality, such as Deep Convolutional GAN (DCGAN) [155] and the Style-based GAN (StyleGAN) [89].

3 Interferometric Particle Imaging for Particle Sizing in Front-, Side- and Backscatter

This chapter addresses RQ1, which explores the extension of IPI from the front-scatter to the backscatter regime. It presents a rigorous framework for conducting IPI across all scattering regimes, including backscatter, beginning with the theoretical foundation and its adaptation to the backscatter configuration. The second part of the chapter examines RQ2, focusing on how uncertainties and limitations in IPI are influenced when applied in side- or backscatter configurations, in comparison to frontscatter.

The results discussed in this chapter build upon the article *Interferometric Particle Imaging for Particle Sizing in the Front-, Side-, and Backscatter Region – Towards Single Optical Access IPI* by Sax *et al.* [SDK25].

3.1 An alternative Transfer Function for Interferometric Particle Imaging

IPI is a powerful technique for sizing small transparent particles within comparatively large fields of view. "In the literature on IPI, the front-scatter region with scattering angles of 66° for water droplets in air [98, 64, 139] and 45° for air bubbles in water [90, 144], is used predominantly due to its superior scattering efficiency and higher signal-to-noise ratio (SNR). More recently, scattering angles of $90 - 110^\circ$ have also been used [222, 170, 42]. However, there is still a lack of research on the use of IPI in the backscatter region. This is probably due to the lower SNR in the backscatter region, which makes the front-scatter region easier to deal with." [SDK25]

To determine particle size in IPI, Eqs. (2.10) and (2.11) serve as the standard transfer functions relating the number of fringes in the PI to the particle size. "The problem with Eqs. (2.10) and (2.11) is that they are Taylor series expansions [176] limited to the forward scattering region (i.e. small θ) and to the scattering orders $p = 0$ and $p = 1$ that dominate this region. Eq.(2.10) is restricted to $\theta < 2 \arccos(1/m)$, which represents the total reflection of the $p = 1$ ray at the surface preventing the light from entering the sphere ($\beta_i^{(1)} \geq 0$). Eq.(2.11) is limited to $\theta < 2 \arccos(m)$ because the $p = 0$ and $p = 1$ glare point collapse into a single point. Consequently, Eq. (2.10) and (2.11) cannot be used in the side- or backscatter regime (i.e. large θ), since the angle relations don't hold true for this range and other scattering orders can dominate these regions." [SDK25] Alternatively, the relationship solely between the glare point positions $w_{p_i}^{\text{phys}}$ and the fringe frequency F of the interference pattern, see Chapter 2.2.3, may be considered, in order to derive an alternative transfer function between the fringe pattern and the particle diameter. "This relationship is independent of the scattering angle and therefore, avoids the first problem of Eqs. (2.10) and (2.11). This is done by introducing the glare point distance between two glare points

$\Delta_{\text{GP}}^{\text{phys}} = |w_{p_i}^{\text{phys}} - w_{p_j}^{\text{phys}}|$, compare Eq. (2.9). With the formulae from Young's fringe experiment, the glare point distance can be determined from [19]

$$1/F = \frac{\lambda_0 d_s / n_2}{\Delta_{\text{GP}}^{\text{phys}}}, \quad (3.1)$$

where d_s is the distance between the particle and the observer (i.e. the camera chip), and λ_0 is the wavelength of the light in vacuum." [SDK25]

The relationship between the glare point distance and the resulting fringe pattern is illustrated in Fig. 3.1. The analogy to Young's double-slit experiment becomes apparent, particularly in the similarity between glare point spacing and slit separation; see also Fig. 2.8b.




Figure 3.1: The relationship between glare point distance and the resulting interference patterns is visualised. As the glare point distance increases, the fringe frequency correspondingly increases. While the angular frequency remains constant along the z -axis, the interference pattern progressively spreads out with increasing distance from the focal plane.

"Eq. (3.1) is constrained to the propagation of light waves in open space. However, to change the direction of light and/or separate different fluids, most measurement scenarios involve optical elements such as e.g. lenses and/or windows, and therefore, a generalization to arbitrary paraxial systems is needed. Such systems can be comprehensively described by ray transfer matrices (also known as ABCD matrices)[22, 60]. [...]

The defocusing parameter B_{tot} replaces the expression d_s/n_2 , for which it can be interpreted as the equivalent distance for propagation through free space representing the propagation through the optical system (with an imaging lens etc.). [...] Shen *et al.*[179] incorporated this approach into Eq. (3.1) to formulate

$$1/F = \left| \frac{\lambda_0 B_{\text{tot}}}{\Delta_{\text{GP}}^{\text{phys}}} \right|. \quad (3.2)$$

Whilst this equation solves the issue of the scattering-angle dependence, a formulation of the relationship between the glare point distance and the particle diameter still needs to be found, in order to perform IPI. As described above, the glare point position and visible glare point orders do change with the scattering angle and the refractive index. For this reason, a general formulation is introduced as

$$\Delta_{\text{GP}}^{\text{phys}(p_i, p_j)}(\theta) = |w_{p_i}^{\text{phys}}(\theta) - w_{p_j}^{\text{phys}}(\theta)| = \frac{d_P}{2} |w^{(p_i)}(\theta) - w^{(p_j)}(\theta)| = \frac{d_P}{2} \Delta_{\text{GP}}^{(p_i, p_j)}(\theta), \quad (3.3)$$

which describes the glare point separation as a function of the particle diameter and the position of any two glare points (p_i, p_j) . The glare point spacing introduces information from the scattering

model into the interference model. Eq. (3.3) can be introduced into Eq. (3.2) to formulate a general relationship between the fringe frequency F and the particle size d_P via

$$d_P = 2 \left| \frac{\lambda B_{tot}}{\Delta_{GP}^{(p_i, p_j)}(\theta)} \right| F. \quad (3.4)$$

Unlike Eqs. (2.10) and (2.11), Eq. (3.4) can be applied to any scattering angle $\theta = [0, 180^\circ]$, but requires prior knowledge of the relative glare point spacing $\Delta_{GP}^{(p_i, p_j)}(\theta)$. The GO model provides information on the position of the glare points but not on their intensity. The intensity I of the glare points plays a crucial role in determining which glare points form the interference pattern since the intensity of individual glare points p_i might vary by several orders of magnitude with the scattering angle. Consequently, the interference pattern is not always the result of the same scattering orders, and some scattering angles do not give rise to visible interference patterns at all. A further model is, therefore, required to determine the glare point spacing that produces the interference pattern." [SDK25]

3.2 Theoretical Model for IPI in Front- Side- and Backscatter

"This section presents a comprehensive method to choose a scattering angle in any scattering regime (front-, side- and backscatter) for IPI. Two main questions are addressed: (i) Which scattering orders form visible interference patterns at which angles and (ii) what is the underlying spacing of the respective glare points that form the interference pattern?"[SDK25]

3.2.1 Visibility of Interference Patterns

"The intensity of the glare point of each scattering order depends on the scattering angle and the refractive index. As a result, interference patterns are not as clearly visible from all angles, if they are visible at all. Therefore, the first step is to identify scattering angles that produce visible interference patterns. In analogy to Young's fringe experiment, where the slits can be considered as point emitters of light, the glare points are modeled as point emitters, compare Fig. 2.8b. Thus it can be postulated that the scattering process at a particle leads to a visible interference pattern in the far field, if the glare points act similar to a double slit, as formulated in Theorem 1.

Theorem 1 *A transparent, approximately spherical particle with refractive index m , which is illuminated with monochromatic and coherent light, produces a visible interference pattern at a scattering angle θ , if two glare points of order p_i and p_j have approximately the same light intensity $I^{(p_i)}(\theta) \approx I^{(p_j)}(\theta)$ (Condition 1) and the two glare points p_i and p_j comprise the majority of the light intensity summed over all glare points $I^{(p_i)}(\theta), I^{(p_j)}(\theta) \gg I^{(p_{k \neq i, j})}(\theta)$ (Condition 2).*

Condition 1 ensures that the destructive interference of the two respective waves leads to intensity values close to zero and consequently results in good contrast between minima and maxima of the resulting fringe pattern. This condition has already been formulated in the literature[42, 222]. However, *Condition 1* is necessary but not sufficient, so *Condition 2* takes into account all other glare points and ensures that the interference pattern is not eclipsed by light emitted by a third even brighter glare point, see Fig. 3.2. When measuring in the front scatter region, no other glare points then $p = (0, 1)$ dominate and the sufficient condition is fulfilled automatically. For the side- and backscatter regime, however, *Condition 2* is needed.




Figure 3.2: "The concept of visibility and the two conditions are visualized in an example at a single scattering angle. Shown is the intensity over the projected particle surface as determined by the LMT (middle) and the decomposition into individual scattering orders with the Debye-Series expansion (left, right). The glare points are the peaks of intensity." [SDK25]. Figure from [SDK25].

Theorem 1 provides the theoretical basis but needs to be formulated explicitly to be of practical use. Therefore, the intensity of the glare points needs to be quantified. A generally accepted model for the intensity of light scattered by particles in the size range relevant to IPI is the Lorentz-Mie-Theory (LMT) [200]. [...] Van de Hulst [200] derived the complex amplitude of the scattered light in the far-field [...] from the LMT. The intensity is then calculated from the transverse electric (TE) S_1 and transverse magnetic (TM) S_2 modes of the wave via

$$I(\theta) = \sqrt{|S_1(\theta)|^2 + |S_2(\theta)|^2}. \quad (3.5)$$

The LMT also considers the polarization of the light, which is a crucial feature for the contrast and intensity of interference patterns, as will be shown later. [...] The problem with the LMT, however, is that it only provides the integral intensity over all glare point orders, whereas for IPI the intensity of individual glare points is of interest. Starting from the LMT, the integral intensity can be decomposed into individual scattering orders with the Debye-Series expansion of the LMT [40]. The calculation of $S_1(\theta)$ and $S_2(\theta)$ by means of the Debye-Series expansion is complex and computationally intensive. An efficient and straightforward calculation of $S_1(\theta)$ and $S_2(\theta)$ is possible by means of publicly available research software such as e.g. *MiePlot* [107], which has been chosen in the present work; see Fig. 3.3.

Suitable scattering angles for IPI can be identified qualitatively from consideration of Fig. 3.3 in combination with Theorem 1 as shown in Figs. 3.3a and 3.3b. For a quantitative evaluation of the visibility, however, an explicit formulation of Theorem 1 is needed.

Condition 1 ensures that the two glare points forming the interference pattern have the same intensity $I^{(p_i)}(\theta)/I^{(p_j)}(\theta) \approx 1$ and poses the necessary condition. Taking into account the phase difference between the two glare points *Condition 1* can be quantified as follows [42, 222].

$$V_{C1}^{(p_i, p_j)}(\theta) = \frac{2\sqrt{I^{(p_i)}(\theta)I^{(p_j)}(\theta)}}{I^{(p_i)}(\theta) + I^{(p_j)}(\theta)} \approx 1 \quad (3.6)$$

Condition 2 ensures that the glare points p_i and p_j that satisfy *Condition 1* are not eclipsed by other brighter glare points p_k , and thus states the sufficient criterion. In accordance with Eq. (3.6), Eq. (3.7) is formulated to approach one when $I^{(p_i)}(\theta)$ and $I^{(p_j)}(\theta)$ comprise all the intensity of




Figure 3.3: "The complex amplitude S_2^2 of each scattering order is plotted over all scattering angles. Depicted are the Debye plots for (a) a water droplet in air ($m = 1.333 + i1.82E-09$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1476.3$) and (b) an air bubble in water ($m = 1/1.333 + i$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1967.9$). The visibility criteria are visualized: Scattering angles satisfying both criteria are marked green (\bullet), whereas angles only matching the first condition but not the second one are marked red (\bullet)."[SDK25]. Figure from [SDK25]

the scattered light. Therefore the intensity of the considered glare point pair is compared to the integral intensity over all scattering orders.

$$V_{C2}^{(p_i, p_j)}(\theta) = \frac{2\sqrt{I^{(p_i)}(\theta)I^{(p_j)}(\theta)}}{\sum_{k=0}^{\infty} I^{(p_k)}(\theta)} = \frac{2\sqrt{I^{(p_i)}(\theta)I^{(p_j)}(\theta)}}{I_{\text{Mie}}(\theta)} \approx 1 \quad (3.7)$$

Since the Debye series expansion approaches the Mie solution for the sum over all scattering orders $k \rightarrow \infty$, the simplification $\sum_{k=0}^{\infty} I^{(p_k)}(\theta) = I_{\text{Mie}}(\theta)$ can be used. This allows for the exact and faster computation of *Condition 2*. Both Eqs. (3.6) and (3.7) are formulated in order to approach one when the underlying condition is satisfied, with $V(\theta) \in [0, 1]$. Therefore, to give a quantitative description of Theorem 1, the two expressions can be multiplied to give

$$V^{(p_i, p_j)}(\theta) = V_{C1}^{(p_i, p_j)}(\theta)V_{C2}^{(p_i, p_j)}(\theta) = \frac{4I^{(p_i)}(\theta)I^{(p_j)}(\theta)}{(I^{(p_i)}(\theta) + I^{(p_j)}(\theta))I_{\text{Mie}}(\theta)} \approx 1 \quad (3.8)$$

while $V(\theta) \in [0, 1]$ is preserved. The visibility of an interference pattern V_1 or V_2 for a TE or TM wave can be computed directly from the complex amplitudes S_1 or S_2 (Figs. 3.3a and 3.3b) to obtain Fig. 3.4. This allows the visibility to be quantified and the underlying scattering orders to be identified, solving the first problem mentioned in the introduction to this section. Eq. (3.8) is valid for any transparent sphere with a homogeneous refractive index. As outlined above, two cases need to be distinguished for droplets ($m > 1$) and bubbles ($m < 1$). The most common examples in IPI are water droplets in air ($m = 1.333$) and air bubbles in water ($m = 1/1.333$), which will therefore be used as examples in this treatise. For other refractive indices, the values presented in this work may change, but the outlined general effects remain the same for $m > 1$ or $m < 1$.[SDK25]

3.2.2 Visibility Plots and Examples

"In this subsection, the visibility of interference patterns at different angles is discussed and the involved phenomena are elaborated. Fig. 3.4 shows the visibility plot for water droplets in air and air bubbles in water with a diameter of $d_p = 250 \mu\text{m}$ and light with $\lambda = 532 \text{ nm}$. High values

of $V_{1,2}$ correspond to better visibility. High visibility ultimately leads to better SNR (allowing for a better distinction between minima and maxima) and, consequently, to more reliable signal processing. However, the SNR also depends on the intensity of the scattered light itself, which is not considered in the visibility.

As expected, Fig. 3.4 predicts good visibility in the front-scatter regime at the angle of 66° for water droplets in TE-mode and at 45° for bubbles in water in TE-mode. It can also be seen that the orders $p = (0, 1)$ form the interference pattern, which is in agreement with the literature [98, 64, 139, 144, 90]. Fig. 3.4 shows that these angles are only preferable in the TE-mode, but in the TM-mode, 84.5° for droplets and 41.6° for bubbles are preferable, while the $p = (0, 1)$ pair remains the dominant one. This demonstrates the importance of the polarization of the light used for IPI. However, for the scope of this work, the less explored side- and backscatter regions are of greater interest.




Figure 3.4: "Visibility over the scattering angle for a water droplet in air ($m = 1.333 + i1.82 \cdot 10^{-9}$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1476.3$) in (a) TE- (V_1) and (b) TM-mode (V_2), and for an air bubble in water ($m = 1/1.333 + i$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1967.9$) in (c) TE- (V_1) and (d) TM-mode (V_2). A sufficient scattering angle for IPI is given whenever the visibility plot has a value close to one. The glare points parings $(0,1)$, $(0,2)$, $(0,3)$ and $(1,2)$ are shown. In addition, the rainbow angles of the $p = 2$ and $p = 3$ scattering orders and Alexander's dark band (ADB) are marked in (a) and (b). The Brewster angle is marked in (d)."[SDK25]. Figure from [SDK25]

For water droplets and TE-mode, see Fig. 3.4a, high visibility is achieved at scattering angles of 111° - 129° with $p = (0, 3)$ and 137.5° - 165° with $p = (0, 2)$. Both regions are separated by Alexander's dark band (ADB) [92], a region in which the intensity of the $p = 2$ and $p = 3$ glare points plummet. ADB is encased by the rainbow angles, a phenomenon in which glare points of

the same order collapse into a single glare point, increasing the intensity [201]. Both regions show a large variability of the visibility. This is undesirable because as a particle moves through the field of view, the effective scattering angle varies slightly (up to a few degrees) and consequently interference patterns will be visible at some positions in the image but not at others. Water droplets in TM-mode, see Fig. 3.4b, achieve similarly high visibility at scattering angles of 111° - 129° with $p = (0, 3)$ and for angles larger than 137° with $p = (0, 2)$, again separated by ADB. A major advantage of the TM-mode is the lower variability of the visibility – especially in the range of 111° - 129° , which renders TM a more reliable and robust choice.

For bubbles in TE-mode, see Fig. 3.4c, the front-scatter region is the only region to yield a visibility of $\approx 100\%$. The $p = (0, 2)$ pair dominates most of the side- and backscatter regions but, with lower visibility. The $p = (0, 3)$ pair reaches up to 86% visibility at 173° - 179° but is again subject to strong fluctuations. Bubbles in TM-mode achieve high visibility at 96° - 101° and 114° - 179° with $p = (0, 2)$, see Fig. 3.4d. The two regions are separated by the Brewster angle, a region where the $p = 0$ ray hits the surface at an angle at which the transmission coefficient is one and the reflection coefficient is zero. Consequently, no light is reflected at this point, compare Fig. 3.3b at 106° . Visibility drops sharply as the intensity ratio at this point is very uneven. The two regions of high visibility show almost no fluctuations. Furthermore, the visibility in the range of 114° - 179° is almost constant, making it ideal for use in IPI. Around 179.8° , the $p = (0, 3)$ pairing reaches up to 85% visibility but is subject to strong fluctuations.

The above elaboration demonstrates that the TM-mode is favorable for the side- and backscatter regimes, for both droplets and bubbles. This is remarkable information, since for the forward-scatter for droplets, usually the TE-mode is used, while for bubbles both modes work in the forward-scattering region. Other effects such as ADB and the Brewster angles lead to regions of the scattering angle in which no interference pattern is visible and IPI cannot be performed. It is important to note that high visibility is necessary for IPI, but for practical use, other effects such as SNR and uncertainties need to be taken into account too, which will be discussed in the following sections."[SDK25]

3.2.3 Determination of the Glare Point Spacing

"The visibility plots contain the information about appropriate scattering angles θ and the glare point paring (p_i, p_j) that forms the interference pattern. From this information, the value of the glare point spacing $\Delta_{GP}^{(p_i, p_j)}(\theta) = |w^{(p_i)}(\theta) - w^{(p_j)}(\theta)|$ can be determined as required for Eq. (3.4). This section presents two ways to compute the glare point position $w^{(p_i)}(\theta)$, which are interrelated but have different advantages. Additional effects not considered in the visibility considerations are discussed.

The first way to compute the glare point positions is the derivation of the intensity over the particle surface from the complex amplitude S ."[SDK25] Eq. (2.39), introduced in Chapter 2.3.1, can be employed to determine the intensities $A_{GP,1}(w)$ and $A_{GP,2}(w)$, corresponding to the scattering amplitudes S_1 and S_2 , respectively, over the projected particle surface w . "Note that while the position w of the glare points do not change with the polarization, the intensities $A_{GP,1}(w)$ and $A_{GP,2}(w)$ of the glare points differ, as contrasted in Figs. 3.5 and 3.6 for droplets and bubbles, respectively.

Eq. (2.39) effectively describes the propagation of the wave through a lens and shows how a camera would see the glare points in focus. The angular range $\Delta\theta = \tan(\alpha_{AP})$ follows directly from geometrical considerations and is defined by the opening angle of the camera α_{AP} set by the aperture, which is the same angle as used in Eqs. (2.11) and (2.10). The opening angle is calculated

from the triangle formed by the effective optical distance of the particle from the lens $z_{P,\text{lens}}/n_2$ and the aperture diameter D_a with $\tan(\alpha_{AP}) = D_a n_2 / (2z_{P,\text{lens}})$. While Van de Hulst [201] introduced Eq. (2.39) for the integral intensity of all glare points (LMT), the complex amplitude $S^{(p_i)}(\theta)$ from the Debye-Series expansion of the LMT can be inserted instead of S_{Mie} to decompose the intensity into the individual scattering orders, compare Fig. 3.2. The result is shown in Figs. 3.5 and 3.6 for droplets and bubbles, respectively, and hereafter referred to as glare point plots. The glare point position of scattering order p_i is then simply determined by finding the peak position of $A_{\text{GP}}^{(p_i)}(w)$. As Eq. (2.39) is a short space Fourier transform, the uncertainty relation $\Delta\theta\Delta w = 1/x_{\text{size}}$ applies[201]. This means that for smaller particles the peaks become less distinct, to the point where the glare point model collapses, i.e. when the particle approaches the size of the wavelength of light. For demonstration purposes, a large particle diameter has been arbitrarily chosen to make the discussed effects more salient.

The glare point plots allow the intensity of the glare points to be assessed in a straightforward manner in addition to finding the glare point spacing, as shown in Fig. 3.5a. This allows to judge the corresponding visibility intuitively. However, the disadvantage of this method is that each plot shows only a single scatter angle. Fig. 3.5 shows a range of scattering angles with high visibility and other interesting effects. The first effect that can be seen, is the influence of polarization on the intensity of the glare points, as can be identified from the comparison of Fig. 3.5a and Fig. 3.5b at $\theta = 66^\circ$.




Figure 3.5: "Intensity over the particle surface (the glare points are the peaks): The scattering intensity $A_{\text{GP}}(w)$ is plotted over the particle surface w for a water droplet in air ($d_P = 250 \mu\text{m}$, $\lambda = 532 \text{ nm}$, $\Rightarrow x_{\text{size}} = 1476.3$ and $m = 1.333 + i1.82 \cdot 10^{-9}$). The scattering uncertainty $\Delta\theta = 3.7^\circ$ was chosen based on the optical system used for the validation experiment. The scattering angle θ_0 around which $A_{\text{GP}}(w)$ was calculated is shown at the top of each plot. The intensities of the scattering orders $p=0$ (●), $p=1$ (●), $p=2$ (●) and $p=3$ (●), as well as the integral intensity of all scattering orders (Mie-intensity) (●) are shown. Unless indicated all plots are on the same intensity scale." [SDK25]. Figure from [SDK25]

One effect that is not considered in the visibility considerations is that the same scattering order can produce more than one glare point, as discussed by Van de Hulst [201], compare Fig. 3.5b $\theta = 160^\circ$. The intra-order split-up of glare points, can cause a significant loss in intensity of the glare point as the light intensity is distributed over more than one glare point, compare Fig. 3.5a $\theta = 116^\circ$ and $\theta = 129^\circ$. This reveals a shortcoming of the visibility consideration, as it only considers the integral intensity of the scattering order at one scattering angle. If a single order splits into

multiple glare points, then this integral intensity is also split between the two glare points of the same order, which cannot be captured by *Condition 1*. Such cases can nonetheless still be used for IPI and even have the advantage of introducing multiple frequencies due to the three glare point spacings, see Fig. 3.5a $\theta = 160^\circ$. Dehaek *et al.* [42], for instance, used a third glare point of a third order in sidescatter to create redundancy due to multiple frequencies being present in the interference pattern. This idea might also be extended to the use of intra-order split-ups to obtain multiple frequencies from only two orders. Combined with the reasonably high intensity in backscatter for droplets, scattering angles around 160° can, therefore, be attractive for IPI. Note that the intra-order split-up for $p = 3$ at $\theta = 116^\circ$ occurs only for the TE-mode but not in the TM-mode, again highlighting the importance of polarization for IPI. At certain angles – the so-called rainbow angles [201] – the two glare points of the same order collapse into a single glare point. The brightness of this rainbow glare point is substantially higher, since all light intensity of that order is comprised of a single glare point, compare Fig. 3.5a $\theta = 129^\circ$ and 137.9° . If two glare points of the same order are undesirable, measuring at the rainbow angles [201] can get around this problem. However, measurement at the rainbow angles is not preferable due to the uncertainty characteristics as will be elaborated further in the following sections.




Figure 3.6: "Intensity over the particle surface (the glare points are the peaks): The scattering intensity ($A_{GP}(w)$) is plotted over the particle surface w for an air bubble in water ($d_P = 250 \mu\text{m}$, $\lambda = 532 \text{ nm}$, $\Rightarrow x_{\text{size}} = 1967.9$ and $m = 1.333 + i1.82 \cdot 10^{-9}$). The scattering uncertainty $\Delta\theta = 3.7$ was chosen based on the optical system used for the validation experiment. The scattering angle θ_0 around which $A_{GP}(w)$ was calculated is shown at the top of each plot. The intensities of the scattering orders $p = 0$ (●), $p = 1$ (●), $p = 2$ (●) and $p = 3$ (●), as well as the integral intensity of all scattering orders (Mie-intensity) (●) are shown. Unless indicated all plots are on the same intensity scale." [SDK25]. Figure from [SDK25]

Bubbles show less complex phenomena, see Fig. 3.6. The TM polarization shows more uniform intensities between the glare points in side and backscatter, compare Fig. 3.6b. However, the TE mode has the advantage of higher intensities, compare Fig. 3.6a $\theta = 99^\circ$ - 154° , which potentially compensates for the lower visibility (due to uneven intensity) with higher intensity-based SNR. For the TE-mode, bubbles also show a third $p = 3$ glare point in the side- and backscatter region, see Fig. 3.6a $\theta = 99^\circ$ - 154° , and a third $p = 2$ glare point at 45° . Again, this third glare point could be used for multi-frequency IPI as suggested by Dehaek *et al.* [42]. Unlike droplets, the side- and backscatter regimes have significantly lower intensities compared to the frontscatter regime, potentially resulting in lower SNR." [SDK25]




Figure 3.7: "Glare point maps: Glare point position w over the scattering angle visualized for the reflected and first three orders of refracted light rays for a droplet (a) and a bubble (b). For the droplet in (a), the rainbow angles and ADB are marked. For the bubble in (b) it is shown how to measure the glare point spacing from the plots."[SDK25]. Figure from [SDK25]

The second method for deriving the glare point spacing is based on the GO model, as illustrated in Fig. 2.7 and 2.8b. This approach employs Eq. (2.13) [201] to determine the angle of the ray relative to the surface tangent. For clarity and ease of reference, Eq. (2.13) is restated below:

$$\theta' = 2 \left(\beta_i^{(p)} - p \arccos\left(\frac{1}{m} \cos(\beta_i^{(p)})\right) \right) = 2\pi k + q\theta. \quad (3.9)$$

"The two unknown constants $q \pm 1$ and $k \in \mathbb{N}$ are needed in order to formulate Eq. (2.13) without loss of generality for $m \leq 1$ and arbitrary p . Solving this equation for $\beta_i^{(p)}$, then allows to compute the glare point position from $w^{(p)} = q \cos(\beta_i^{(p)})$. However, Eq. (2.13) is not easy to solve analytically and it is necessary to find values for q and k for which the equation has a real solution within $\beta_i^{(p)} \in [0, 90^\circ]$. Therefore, an algorithm (Alg. 1) for the indirect solution of Eq. (2.13) and the determination of the values of q and k is proposed and provided in Appendix A.2. Furthermore, Tab. A.1 lists values of q and k for $m \leq 1$ and $p \in [0, 3]$ (see also Appendix A.2). With the solution of Eq. (2.13), the glare point position can be plotted over the scattering angle, see Figs. 3.7a and 3.7b, hereafter referred to as glare point maps. The glare point separation is determined by the horizontal distance between the two lines of the desired order, see Fig. 3.7b. Unlike the first approach, no intensity information can be derived and therefore shortcomings of the visibility method cannot be identified with this approach. Accordingly, this also means that polarization effects cannot be assessed. However, the glare point separation between a range of scatter angles can be easily compared using these glare point maps. This allows to investigate the change of glare point spacing with the scattering angle. This is a crucial tool for deciding, which scattering angle to choose since the GP spacing has a major influence on the uncertainty and limitations of IPI (see next sections). Furthermore, the rainbow angles can be easily identified from this approach (condition: $d\theta/dw = 0$; i.e. local min. or max.[201]).

The relationship between the two approaches can be understood as the three-dimensional interrelation of intensity, scattering angle, and glare point position, which is accordingly visualized in Fig. 3.8. Both intensity and polarization, as well as the change in the glare point spacing over scattering angles, are crucial quantities to consider for IPI. It is therefore advisable to consider both glare point plots (Figs. 3.5 and 3.6) and glare point maps (Fig. 3.7) when choosing a scattering angle to take all effects into account."[SDK25]




Figure 3.8: "Interrelationship between intensity, scattering angle θ and glare point position w . Glare point plots represent vertical slices placed on the glare point maps (horizontal slice)."[SDK25]. Figure from [SDK25]

3.3 Experimental Validation of Backscatter IPI

"In this section, the methodology derived above is validated experimentally. Eq. (3.1) is the formulation for a double slit in Young's fringe experiment and has been well validated in physics. Also Eq. (3.2) introduced by Shen *et al.*[179] has been validated by e.g. Zhang *et al.*[222] by measuring the same bubble size distribution of air bubbles in oil at scattering angles of 45° and 90° . Eq. (3.4) is a rearrangement of Eq. (3.2) with a model for the glare point position added in Δ_{GP}^{phys} . To test the methodology outlined in this"[SDK25] thesis, "a bubble generator with nominally constant bubble size distribution is chosen and characterized with IPI. Given the distribution function of the generated bubble sizes, i.e. its probability density function (PDF), can be considered constant and, consequently, the measured bubble size distributions should always be the same – regardless of the applied scattering angle. The test of the above working hypothesis is as follows: If IPI measurements based on the methodology outlined above all result in the same PDF, and interference patterns are visible at the predicted scattering angles, then the suggested method is valid. Based on Fig. 3.4c, six different scattering angles were chosen, i.e. 45° , 99° , 116° , 122° , 135° , and 154° inside the water at the bubble. To validate Eq. (3.2)[179] the results at an angle of 45° are additionally compared with Eq. (2.11)[144, 90]."[SDK25]

3.3.1 Experimental Procedure

"In order to conduct the outlined testification approach, an experimental setup as shown in Fig. 3.9 was constructed that was able to generate a constant distribution of bubbles. The setup allowed the scattering angle to be varied without changing the optical path since even small changes in the optical path can affect the result, see Fig. 3.9a. To achieve a bubble generation with a constant size distribution, a generator based on electrolysis was constructed. The bubble generator consisted of two stainless steel rods sheathed in black rubber for electrical insulation except at the ends. Both steel rods were connected to a power source, which provided a constant current and voltage throughout the measurements. A sheathed silver wire was connected to the end of the cathode (via a borehole). The wire was bent into a hook shape and only the last 0.5 mm of the wire was left unsheathed, compare Fig. 3.9c, right. In this way, bubbles were generated only at the tip of the hook (known location). The field of view was placed directly on the top of the hook to ensure that the measured bubbles were at a known distance from the camera chip. The medium surrounding the bubbles was distilled water with NaCl added to ensure conductivity. The generated hydrogen bubbles had a relative refractive index of $m = 1/1.333$. The point of bubble generation was placed

in a water tank (glass $n = 1.52$, $20 \times 20 \times 20 \text{ cm}^3$, 4 mm wall thickness) and fixed at a submersion depth of 4.1 cm.




Figure 3.9: "The optical system of the validation experiment is shown. Depicted are a schematic principle of the experimental setup (a) and an image of the real experimental setup (b). A schematic close-up of the bubble generator (cathode) and the field of view (FOV) relative to the tip of the cathode-hook are shown (c). The physical model for the optical path from the bubble to the camera is shown up in a close-up (c)"[SDK25]. Modified figure from [SDK25]

A camera (PCO Pixelfly: CCD chip, 14 bit, $1392 \times 1040 \text{ px}$, pixel size $6.45 \times 6.45 \mu\text{m}^2$ equipped with a Nikon Micro-Nikkor imaging lens: effective focal length 105 mm, aperture number $f_{\#} = 4$) was positioned on an arm outside of the tank at a fixed distance from the glass. A laser (Quintel Evergreen Nd:YAG $\lambda = 532 \text{ nm}$, 200 mJ/pulse) was positioned on a rotating arm to allow for different scattering angles without a change in the optical path from the bubble to the camera, compare Fig. 3.9b. Note that the laser head was oriented such that measurements took advantage of the favorable TE-mode with correspondingly higher intensity of the scattered light (at the expense of lower visibility), see Fig. 3.6. The laser and camera were synchronized in a single pulse single exposure setup. The optical system is shown in Fig. 3.9c. The distance from the bubble to the inner glass wall was $z_w = 125.5 \text{ mm}$ and the thickness of the glass wall was $z_g = 4 \text{ mm}$ ($n_g = 1.52$). The principal plane was located at the defocused position of $z_a^{\text{defoc}} = 11.71 \text{ mm}$ and the image distance from the principal plane to the camera chip was $z_{im} = 21.81 \text{ mm}$. Since a multi-lens camera lens was used, the position of the principal plane was unknown, so z_{im} and z_a were calculated from imaging a calibration target at the position of the bubble generator. The thin lens equation

$$1/z_{\text{obj}}^{\text{eff}} + 1/z_{\text{im}}^{\text{eff}} - 1/f_{\text{lens}} = 0 \quad (3.10)$$

with $z_i^{\text{eff}} = z_i/n_i$ for the effective distances in combination with the magnification condition provide two equations to calculate z_{im} and z_a^{foc} ."[SDK25] The effective object side distance is $z_{\text{obj}}^{\text{eff}} = z_a/n_a + z_g/n_g + z_w/n_w$. "When moving the camera out of focus by Δz_a , the value $z_a^{\text{defoc}} = z_a^{\text{foc}} + \Delta z_a$

can be calculated (see Appendix A.1 for more detail). The resulting defocusing parameter from the ray-transfer matrix, describing the optical system, was $B_{\text{tot}} = -7.57 \text{ mm}$, see Eq. (A.4). [SDK25] This kind of calibration method is described in Chapter 2.2.2.

3.3.2 Data Processing

A total of 1200 images was recorded per angle and the laser brightness was adjusted at each angle to obtain sufficient intensity. Example results are shown in Figure 3.10 for all investigated scattering angles. The images were pre-processed using mean image subtraction to subtract the background with only little loss of SNR. Morphological closing (dilation and erosion) was applied to the mean-subtracted images followed by a Hough transform to identify the circular particle images (PI). The detections from the Hough transform were transferred into a bounding box around the particle (kernel). This was done for all angles except 154° since the Hough transform had a large detection missrate due to the lower SNR and the bounding boxes had to be drawn manually. The underlying problem for this seeming shortcoming is that the light scattered by larger particles is brighter and therefore larger particles have higher SNRs. As the smaller particles accordingly have lower SNRs, the detection algorithm is unable to detect them and a sampling bias towards larger particles is introduced, thus falsifying the estimated PDF. This limitation necessitated manual labeling of the 154° images to avoid the sampling bias. At the other angles, the PIs were sufficiently distinct from the background for the detection algorithm to operate without bias.

The validation experiment shows that interference patterns are visible for all angles predicted with Fig. 3.4d, as can be seen in Fig. 3.10a. This shows that the visibility criterion formulated in Eq. (3.8) can be used for real experiments to identify applicable scattering angles for IPI. Furthermore, Fig. 3.6a predicts a third glare point at 45° for the TE mode. The effect of a third glare point can also be observed in the experiments, see Fig. 3.10a 45° , as the maxima from the primary glare point pairing also have stripes from a secondary frequency, which is the effect of a third glare point.

A Laplacian filter was applied to the kernel to enhance the edges while preserving their position. The kernel was then averaged along the stripes to increase the SNR. The signal was low-pass filtered using an FFT with a frequency cut-off to further reduce noise. The fringe frequency was then determined in two steps: A peak-finding algorithm was applied to the filtered signal to determine a frequency prior to the peak-to-peak distance. This step was necessary to avoid the subsequent fitting algorithm to fit the envelope frequency (the rectangle function resulting from the defocused circle).

A first-order Fourier function was then fitted on the raw signal (average along the stripes, cp. also Fig. 3.11a alongside uncertainty discussion of Section 3.4.1), starting from the frequency prior. To check the measured distributions for statistical significance, the statistical measure $\sum_{i=3}^N \sigma(d_1, \dots, d_N) / \mu(d_1, \dots, d_N)$ (standard deviation normalized by mean) was examined starting from $i = 3$ sampled bubbles (3 as the minimum for the standard deviation) and gradually increasing the considered sample size up the amount of measured bubbles N . Statistical convergence was observed around $N_{\text{bubbles}} \approx 800$ bubbles. The measurements include at least twice as many particles, except for the 154° evaluation, which was limited to 1000 particles due to manual labeling.

Fig. 3.10b shows that the measured PDFs of the bubble diameter collapse, and according to the testing hypothesis this consequently means that the method works for IPI in any scattering regime. The same size distribution was measured at each scattering angle, which demonstrates that the shift in glare point spacing $\Delta_{\text{GP}}^{(p_i, p_j)}(\theta)$ for Eq. (3.4) was correctly compensated when using Fig. 3.6 and 3.7. Small variations in the PDFs were within the measurement uncertainty, which is quantified in the next section. A closer look at Fig. 3.10a shows that although interference patterns are clearly




Figure 3.10: "Experimental results of the validation experiment. Raw images from the validation experiment are shown in (a). For the image at 45°, a secondary frequency can be observed, which is the result of a third glare point ($p = 2$) in the TE-mode. It can be seen that the particles visually become less distinct from the background and the intensity decreases in the backscatter region. The determined distribution functions for all investigated scattering angles 45°, 99°, 116°, 122°, 135° and 154° (angle in water at the bubble) are plotted in (b). The number of particles in each PDF is shown above the graph. The scattering angle of 45° the PDF was evaluated with Eqs. (2.11) ($45^\circ S$) and (3.4) (45°). The smallest measurable bubble diameter for $\theta = 45^\circ$ was reached and is marked with $d_{p,\min}^{45^\circ}$. The PI brightness in the raw images (a) is lower at higher scattering angles, as the overall intensity of the interference pattern is reduced. This can be observed by comparing the intensity in Fig.,3.12."[SDK25]. Modified figure from [SDK25]

visible, PIs in the backscatter region are significantly less distinct from the background than PIs in the frontscatter region. This shows that although the visibility provides the theoretical basis for finding interference patterns, other aspects such as the SNR need to be considered additionally in real experiments. The reduced signal quality in the backscatter regime and the different values of $\Delta_{GP}^{(p_i, p_j)}(\theta)$ affect the uncertainties and boundaries of the experiment as will be elaborated further in Section 3.4."[SDK25]

3.4 Uncertainties and Limitations of IPI

"This section provides an overview of the main sources of uncertainty in the first section and a further discussion of the influencing factors in the second section."[SDK25]

3.4.1 Sources of Uncertainty

"In the calculation of the particle diameter by means of Eq. (3.4), three terms in the equation are subject to uncertainty: B_{tot} which describes the distance from the focal plane, the glare point spacing $\Delta_{GP}^{(p_i, p_j)}(\theta)$, and the fringe frequency F (assuming λ is well known, which holds true for typical lasers). For the present case the out-of-focus location of the bubbles was quasi-constant, and accordingly considered known and fixed. This means that B_{tot} can be obtained from the calibration of the measurement. The camera was moved on a precision traverse to focus on a calibration target so as to obtain the unknowns z_{im} and z_a as described in the Section 3.3. The precision of this calibration step is limited by the depth of field of the mounted lens, which was experimentally determined to be less than 10 μm . This uncertainty can be incorporated into the thin lens equation and the magnification condition to investigate the effect on the calculation of z_{im} and z_a , and the subsequent effect on B_{tot} . For the given depth of field and optical system the resulting uncertainty of $\Delta B_{\text{tot}} = 0.005\%$ can be considered negligible. Note, however, that in a

scenario where the distance from the particle to the camera is not fixed, B_{tot} becomes a function of the distance and its uncertainty is likely to be significantly larger." [SDK25]. This issue is discussed in more detail in Chapter 4.

"The uncertainty $\Delta_{\text{GP}}^{(p_i, p_j)}(\theta)$ of the second term in Eq. (3.4) stems from the uncertainty in the scattering angle. This uncertainty can arise either from the measurement setup (positioning of laser and camera) or from the motion of the particles perpendicular to the stripes of their interference pattern (here left and right in the image $\theta = \theta_0 \pm \Delta\theta$). In both cases, this is a deterministic error. As shown in Fig. 3.7 even such small changes in θ can change the glare point spacing. This relationship is highly non-linear and depends on the scattering angle. A further discussion and quantification of this uncertainty is given in Section 3.4.2. For the present validation experiment, a conservative upper limit of $\Delta\theta$ of 1° was chosen. The resulting new glare point spacing $\Delta_{\text{GP}}^{\text{Err}}$ can be read directly from Fig. 3.7. For the scattering angles used, the resulting uncertainties are given in Tab 3.1. A different glare point spacing changes the frequency of the interference pattern, thus mimicking a change in particle diameter and corresponding changes in the interference pattern (both change the glare point spacing). This, in turn, also explains the variation of observed fringe frequencies for similar particles at different scattering angles.

Table 3.1: "Uncertainty as a systematic error introduced by an angle uncertainty of 1° " [SDK25]. Table from [SDK25]

θ	45°	99°	116°	122°	135°	154°
$ \frac{\Delta_{\text{GP}}^{\text{Err}} - \Delta_{\text{GP}}}{\Delta_{\text{GP}}} $ [%]	3.40	1.07	1.42	1.63	2.17	3.78

The last term to consider is the fringe frequency, which is subject to measurement uncertainty and occurring stochastic error. The frequency measurement is affected by a variety of factors such as SNR, optical aberrations, and the performance of the algorithm used to extract the frequency, all of which render a theoretical quantification difficult. Therefore, an experimental approach was used to estimate an upper bound for the uncertainty. The algorithm to extract the fringe frequency was applied four times to the same interference pattern, each time using a different kernel, see Fig. 3.11a (upper part). Each kernel was averaged along the stripes to obtain a signal, which would result in slightly different frequencies (compare Fig. 3.11a, lower part). The absolute spread of the frequencies determined would then give an estimate of the frequency uncertainty. For further investigation, the peak-to-peak distance in the fringe signal $D_{\text{pp}} = 1/F$ is used as the complement metric of F for more intuitive handling. This measure describes the distance of two maxima in the interference pattern, see Fig. 3.11a (lower part). The algorithm used showed a spread between 0.1-0.9 px in D_{pp} for the same particle and is, therefore, sub-pixel accurate in the frequency determination. However, a conservative upper limit of $\Delta D_{\text{pp}} = 1$ px was chosen for further investigation.

This metric then allows to quantify some fundamental properties of the frequency measurement. The lowest frequency that can be measured (i.e. the smallest particle) corresponds to just one wavelength present in the PI, i.e. $D_{\text{pp}} = d_{\text{PI}}$, where d_{PI} is the PI diameter. The upper limit is given by the Nyquist criterion, for which the distance between peaks is $D_{\text{pp}} = 2$ px, which is the minimum distance between two maxima to be distinguishable without aliasing effects (giving the largest measurable particle). The relationship between particle diameter and peak-to-peak distance is non-linear $D_{\text{pp}} \sim 1/d_{\text{P}}$ (since $F \sim d_{\text{P}}$), see Fig. 3.11b. For small particles, the frequency is highly resolved (here 130 px per signal wavelength), whereas for larger particles the frequency is less well




Figure 3.11: "Figure (a) shows the different kernels used to average the signal along the maxima of the interference patterns and the resulting signal across the PI. The intensity in (a)(top) is inverted for better visibility. The peak-to-peak distance between the maxima of the interference pattern over the radius ($\theta = 45^\circ$) is shown in (b). Figure (c) shows the relative error (for an uncertainty of 1px) in the measurement of the peak-to-peak distance over the particle diameter for different scattering angles (for the setup used in the validation experiment). In addition, the upper and lower limits of the measurable particle diameters at each scattering angle can be determined from the respective ends of the error lines."[SDK25]. Figure from [SDK25]

resolved. This is directly related to the measurement uncertainty of the diameter. As shown in Fig. 3.11b, Δd_p depends on the gradient $|dD_{pp}/dd_p|$ and is approximated as

$$\frac{dD_{pp}}{dd_p} \approx \frac{\Delta D_{pp}}{\Delta d_p}. \quad (3.11)$$

Eq. (3.11) can be inserted in Eq. (3.2) to obtain

$$\Delta d_p(d_p, \theta) \leq \left| \frac{\lambda B_{tot} \Delta D_{pp}}{\Delta_{GP}(\theta)} \right| d_p^2 \quad (3.12)$$

which depends on the particle diameter and the scattering angle due to the glare point spacing. This relation can be used to investigate the relative error over the particle diameter, for different scattering angles, see Fig. 3.11c. Note that Eq. (3.12) gives an estimate for the upper limit, but doesn't quantify the uncertainty exactly. It can be seen from Fig. 3.11c that the scattering angle changes the uncertainty of the frequency determination for a constant particle size. This means that the upper and lower limits of measurable particle sizes are also dependent on the scattering angle. The corresponding uncertainty is, therefore, shifted in the same way. For bubbles, the largest diameter may be measured in frontscatter, while the smallest bubbles may be measured in sidescatter, with the backscatter regime approaching the uncertainty of the frontscatter regime. This is a crucial insight, as an experimentalist might then choose a scattering angle based on the expected particle size distribution. This also shows that a scattering angle should be chosen to measure close to the lower limit, as the uncertainties are significantly lower. The frequency uncertainty approaches 50% at the Nyquist diameter and is therefore the dominant uncertainty for large particles."[SDK25]

3.4.2 Main Influences on the Uncertainty

"The last section discussed the quantities that are subject to uncertainty, this section focuses on the main sources of these uncertainties. While Figure 3.11c would suggest very similar properties between the front- and backscatter regime for bubbles, this is not the case due to the significantly lower SNR in backscatter. The SNR is a result of the PI intensity compared to the background noise on the camera chip. The performance of the algorithm to measure ΔD_{pp} is dependent on the SNR. While this relationship depends on the algorithm, it can be generally assumed that better

performance (i.e. lower ΔD_{pp}) is obtained with higher SNR. The SNR depends on several factors, mainly the particle diameter, the scattering angle, and the amount of defocusing z_P . Consequently, Eq. (3.12) must be modified to

$$\Delta d_P(d_P, \theta, z_P) \leq \left| \frac{\lambda B_{\text{tot}}(z_P) \Delta D_{pp}(\text{SNR}(\theta, z_P, d_P))}{\Delta_{GP}(\theta)} \right| d_P^2 \quad (3.13)$$

which accounts for the fact that ΔD_{pp} is not constant. One important factor is the particle size, since larger particles result in brighter interference patterns, as can be proven by the LMT [200]. This suggests that ΔD_{pp} becomes smaller with larger particles, counteracting $\Delta d_P \sim d_P^2$. Furthermore, this means that Eq. (3.12) overestimates the uncertainty for large particles and underestimates it for smaller ones. This is advantageous, as it indicates that large particles can be measured with smaller uncertainties than Eq. (3.12) suggests." [SDK25]

The intensity increases with the particle diameter, suggesting better SNR for large particles. "A second influence on the SNR is the distance of the particle from the focal plane z_P , which is also taken into account in B_{tot} . The intensity of the PI decreases with the distance from the focal plane, resulting in a lower SNR. As the diameter of the PI increases approx. linearly with the distance from the focal plane [57], the area of the PI increases by the power of 2. As the light intensity is distributed over a larger PI area, the intensity of the PI becomes accordingly weaker and the SNR decreases. At the same time, however, a more defocused particle results in larger PIs [179, 57] as indicated in Fig. 2.8b, which in turn results in better interference pattern resolution and lower ΔD_{pp} . Consequently, there are two counteracting effects on ΔD_{pp} when z_P is changed: The increase in the resolution of the interfering pattern is countered by the decrease in the SNR of the pattern. The increase in fringe frequency resolution is linear with sufficient distance from the focal plane [146], while the intensity loss decreases with the power of 2. The linear increase of resolution coupled with the quadratic decrease of SNR indicates that there is an optimal amount of defocusing for the PI, which is in agreement with earlier PI-characterization studies [116].

Finally, the intensity of the glare points depends on the scattering angle, so the SNR is also a function of θ . The relationship between scatter angle and intensity is highly non-linear. To assess the SNR of the interference pattern, the intensity of its maxima can be examined. These can be calculated from the combined intensity of the two glare points forming the pattern $I_{PI, \text{max}} = |S^{(p_i)} + S^{(p_j)}|$, see Fig. 3.12. This quantity provides a brightness measure for the maxima of the interference patterns (i.e. its distinction from the background - SNR), while the intensity difference between minima and maxima of the pattern is evaluated on the grounds of the visibility. Consequently, both should be taken into account, as both influence the signal quality. For droplets in TE- and TM-mode a strong decrease of intensity from the front- to the side scatter can be seen, compare Fig. 3.12a and 3.12b. From the sidescatter regime towards the $p = (0, 3)$ rainbow angle the intensity increases with the scattering angle, while from the $p = (0, 2)$ rainbow angle the intensity decreases towards even larger scattering angles. As observed for the visibility, the intensity is also low in the ADB, making this angle range unfavorable for IPI. In general, the backscatter regime for droplets shows higher intensities, making it more favorable than the side scatter regime. However, large fluctuations in intensity, see e.g. Fig. 3.12b, cause large variations in SNR for even small changes in angle, as might occur for particles moving through the field of view.

For bubbles also a strong decline of intensity is observed when moving beyond the frontscatter regime, compare Fig. 3.12c and 3.12d. For bubbles in TE-mode the intensity – and therefore the SNR – decreases with increasing scattering angles, which makes experimentation with these angle ranges difficult. Notably, there is a strong increase of intensity for scattering angles close to 180° , while the visibility in this region is also sufficient. For the TM-mode, the intensity decreases




Figure 3.12: "Peak intensity of the interference pattern for a water droplet in air ($m = 1.333 + i1.82 \cdot 10^{-9}$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1476.3$) in (a) TE- and (b) TM-mode, and an air bubble in water ($m = 1/1.333 + i0$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1967.9$) in (c) TE- and (d) TM-mode." [SDK25]. Figure from [SDK25]

towards the Brewster angle, but increases afterwards again, compare Fig. 3.12d. This is valuable information as it indicates an increasing SNR in backscatter with increasing scattering angles. This means that larger scattering angles are potentially favorable for backscatter in TM-mode. The comparison of Figs. 3.12c and 3.12d furthermore reveals that the TE-mode has in general higher intensities in the side- and backscatter regimes. Similar to the TE-mode, the TM-mode shows a sudden increase in intensity for scattering angles close to 180° . However, this higher intensity is within a very narrow angular range, which again causes problems for particles moving through the field of view. Scattering angles approaching 180° present an additional challenge, as coherent backscatter in this range can distort the fringe pattern and must be taken into account. However, since such extreme angles are primarily relevant to microscopic imaging, they are not considered further in this study.

Comparison of Figs. 3.4 and 3.12, in contrast, outlines the complex relationship between visibility and SNR, but also indicates both trade-off and synergetic angle margins, with synergetic angles being preferred.

To test the theoretical approach to SNR issues based on the intensities calculated from the LMT, the PIs from the validation experiment were investigated for their SNR, see Fig. 3.10. While it is difficult to determine the exact SNR, the overall ratio of PI intensity to background intensity $I_{\text{PI}}/I_{\text{BG}}$ can be used as a measure of SNR. The PI intensity is defined here as the average value of the 50 brightest pixels within the bounding box that encloses the PI. The background intensity is defined as the average intensity of a 200×200 px patch in the image that does not contain a PI. A

set of example images as used to compute the SNR is shown in Fig. 3.10a. It is important to note in this context that the brightness of the light source was increased during the validation experiments in backscatter orientation to combat lower overall intensities. Since this in turn likewise increased the background intensity the above-defined SNR estimation provides reasonable information on the relative PI brightness despite varying absolute intensity, as shown in Fig. 3.13. Nevertheless, the experimental data of Fig. 3.13 show a good agreement with the theoretical considerations of Fig. 3.12c, where the intensity decreases strongly from scattering angles of 45° compared to angles in the range of $99\text{--}122^\circ$. The decrease from 122° to 154° is then much smaller compared to the decrease from the front- to the sidescatter regime. The same trend can be seen in the SNR of the experimental data as indicated in Fig. 3.13. This suggests that the combined intensity of two glare points provides a good qualitative guide to the SNR as a function of the scattering angle. However, it becomes clear from both theoretical considerations and experimental data that backscatter IPI requires methods to deal with the significantly lower SNR, such as more robust fringe frequency extraction algorithms. The lower intensities in turn also require better detection algorithms, as smaller particles generate dimmer PIs, which can fall below a given detection threshold upon data processing and accordingly bias any subsequent evaluation of particle-size distribution evaluations. In this context, machine-learning based approaches for flow evaluation and particularly for particle detection have considerably matured in recent years [46, 14] and seem promising for this unavoidable image-processing challenge.




Figure 3.13: "The signal strength (mean intensity of the maxima in every PI) over the background intensity for the images used in the validation experiment. The values correspond to the images shown in Fig. 3.10a. It can be seen that the particles gradually become less distinct from the background and the intensity decreases from the front- over the side- into the backscatter region." [SDK25]. Figure from [SDK25]

Finally, the glare point spacing also plays an important role in Eq. (3.13) as it is the only denominator and depends on the scattering angle. As described earlier, a change in the scattering angle also introduces a change in the fringe frequency for the same particle size, which must be taken into account. Note however that the actual influence on the fringe-pattern changes also depends on the scattering angle itself. The glare point spacing can be obtained from Fig. 3.7, which is plotted over the scattering angle in Fig. 3.14 to emphasize this interplay. There are two factors to consider with respect to the scattering angle – the spacing Δ_{GP} itself and its gradient $d\Delta_{\text{GP}}/d\theta$ upon changes of the scattering angle θ . The former directly scales with the frequency and shifts the uncertainty as well as the lower and upper limits of IPI, as shown in Fig. 3.11c. The latter describes the sensitivity of the glare point spacing to angular changes, which is described by $\text{Se} = |d\Delta_{\text{GP}}/d\theta|$. This sensitivity is an important quantity for typical IPI measurements in a field of view (FOV), because particles likely move through the FOV and accordingly experience different scattering angles in consecutive frames. This means that the glare point spacing changes

– independent from any physical particle-diameter changes – and so does also the uncertainty. It appears, therefore, also beneficial to furthermore quantify the influence of this change in scattering angle for the given experimental setup in order to evaluate the influence of the range of covered scattering angles within the FOV under investigation.




Figure 3.14: "Glare point spacing $\Delta_{GP}(\theta)$ over the scattering angle for droplets ($m = 1.333$) (a) and bubbles ($m = 1/1.333$) (b). The lines are colour coded to indicate the glare point spacing. Only the glare point spacing of the glare point pairing dominant at the scattering angle is shown." [SDK25]. Figure from [SDK25]

For droplets, Fig. 3.14a reveals two important phenomena in the diagrams: around $\theta = 55^\circ$ there is a plateau where the sensitivity Se is zero, which indicates that the influence of the scattering angle on the glare point spacing essentially vanishes. The opposite can be observed for the rainbow angles, where the gradient Se approaches infinity. In these two small indicated angular ranges the sensitivity becomes large and even small location changes inside the FOV or likewise small uncertainties in the experimental setup of the angle will inevitably cause large differences in the frequency of the PI for the same particle diameter (i.e. large uncertainties). This illustrates why rainbow angles, despite their favorable high intensity, are not well suited to IPI. In the range of the intra-order split-ups of $p = 2$ and $p = 3$, where three glare points contribute to the interference pattern, several different gradients and therefore different uncertainties have to be considered at the same time. This shows that a three glare point IPI as proposed by Dehaeck & Beeck[42] provides greater reliability due to redundancy, but also requires a more considerate estimation of uncertainty.

The glare point spacing for bubbles is shown in Fig. 3.14b. Unlike the distinct droplet characteristics, in case of bubbles a similar trend in the Δ_{GP} -slope can be identified for both the front- and the backscatter regime. Particularly for the backscatter region at $\theta > 90^\circ$ the diagram has a near-constant slope, which indicates a constant sensitivity Se . This information is encouraging news for any IPI application in back scatter as it demonstrates that the overall low SNR for such applications does not encounter additional uncertainty estimation issues from angle sensitivities within a chosen FOV." [SDK25]

3.5 Concluding Remarks

In the course of this chapter "it has been shown that IPI can be performed in the backscatter regime. In contrast to the frontscatter regime, interference patterns in the side- and backscatter regimes are formed by different scattering orders and glare point (GP) spacings, which requires

a different transfer function between fringe frequency and particle diameter to account for these changes." [SDK25]

The visibility plots presented in this chapter "can be used to identify suitable scattering angles for IPI as well as to identify the scattering orders that form the interference pattern. The visibility plots have been shown to predict the visibility of patterns from experimental images appropriately. With the GP paring identified, GP-maps and GP-plots can be used to determine the respective GP spacing, as required for the transfer function between fringe pattern and particle diameter. GP-maps have the advantage of giving an overview of the GP spacing at any scattering angle, while GP-plots give more insight into the intensity distribution of the GP and can uncover shortcomings of the visibility approach, e.g. caused by intra-order split-ups.

The study of visibility, intensity, and GP spacing provides several insights that should be considered when selecting a scattering angle. Generally, a high SNR is desirable for a good measurement. For IPI, the SNR depends on both the visibility of the pattern (i.e., the ability to distinguish between minima and maxima in the pattern) and the intensity of the scattered light (i.e., the ability to distinguish the PI from the background). Both factors must be considered together to achieve sufficient SNR, especially in the backscatter regime. In particular, large variations in visibility or intensity can cause difficulties, since measurements with large FOVs include a range of effective scattering angles, which in turn might lead to unfavorable angles for the PIs in certain regions of the recorded image. Consequently, the measurement uncertainty can change with the in-plane position of the particle for scattering angles with strong fluctuations. To manage such intensity fluctuations, a reliable background subtraction, achieved, e.g. through recorded background images or adaptive filtering, becomes crucial for evaluating PIs despite changing SNR. Furthermore, the sensitive change of the GP position for even small angle changes (especially in backscatter) moreover means that the transfer function between frequency and particle diameter effectively becomes a function of the in-plane position of the particle as well.

For droplets, the TE-mode is best suited for measurements in the frontscatter region, but for the side- and backscatter region the polarization should be changed to TM-mode to avoid strong intensity fluctuations and obtain more equal GP intensities. While droplets have the advantage of high intensities in the backscatter regimes compared to bubbles, measurements with droplets in side- and backscatter present several other challenges. The ADB introduces a range of scattering angles at which IPI cannot be performed due to the lack of visible GPs. Furthermore, outside of ADB, the GPs of $p = 2$ and $p = 3$ face an intra-order split, which can be useful for redundancy due to the presence of three frequencies in the interference pattern. However, this effect also requires consideration of three different uncertainties and sensitivities to angle changes, thus rendering uncertainty quantification cumbersome. At the rainbow angles the intra-order split can be avoided, but the sensitivity of the GP spacing becomes large and, moreover, small changes in the in-plane position of the particle might lead to drastic changes in the interference pattern for the same particle size. Therefore, the transfer function becomes highly dependent on the in-plane position of the particle, adding another challenge to the measurement.

Bubbles behave significantly different in the side- and backscatter regimes. For bubbles, both TE- and TM-modes are usable in the side and backscatter regime, with the former giving higher intensities and the latter resulting in better visibility. Interestingly, a region in the side scatter cannot be used for IPI due to the effect of the Brewster angle. Moreover, the fluctuations in intensity and visibility are much smaller for bubbles than for droplets, which renders experimentation with larger FOVs possible without the aforementioned additional uncertainty introduced by the in-plane position. Also, the angle sensitivity of the transfer function is comparable between front and back

scatter. For large scattering angles close to 180° , bubbles allow measurements with high visibility and intensity, but coherent backscatter can disturb the interference pattern.

Three main sources of uncertainty have been identified in the application of IPI. The out-of-plane position of the particle affects the defocusing parameter. This relationship is typically linear, so that uncertainties in the out-of-plane position linearly affect the size measurement. For quasi-2D measurements with known out-of-plane positions of the particles, this error is usually negligible. The second source of uncertainty comes from the GP spacing, which defines the transfer function between the interference pattern and the particle size. The spacing changes the corresponding frequency of the pattern for a given particle size. Consequently, the lower and upper limits of measurable particle sizes can be manipulated with the scattering angle. Since this also affects the resolution per stripe, the accuracy of the frequency determination can also be manipulated by changing the scattering angle.

Finally, the GP spacing affects the aforementioned sensitivity of the transfer function to angular changes, i.e. the in-plane position of the particle. The last influencing factor is the uncertainty of the frequency determination itself, which is mainly determined by the resolution of the fringes and their SNR. The SNR is influenced by the particle size (larger particles have a higher scattering efficiency) and the amount of defocusing (intensity loss due to defocusing), in addition to the visibility and intensity dictated by the scattering angle. The resolution of the PI is determined by the amount of defocusing and the GP spacing, and thus by the particle size and scattering angle. These interrelationships lead to the following two main trade-offs: First, for a given particle size, the scattering angle changes the fringe frequency, which vice versa can be used to calibrate the desired number of fringes for the investigated particle size distributions far away from the Nyquist criterion. This adjustment, inevitably also affects both the visibility and intensity and therefore the SNR of the PI will also change. As such the trade-off is found between a good fringe detection accuracy and a good SNR. Second, more defocused PIs have better spatial resolution of the PI, but at the cost of lower SNR due to intensity loss. Therefore, there is both an optimal particle size and an optimal amount of defocus to minimize the uncertainty in the frequency determination, both of which immediately rely on the given experimental situation and corresponding equipment.

Despite the lower signal quality and more complex handling of uncertainty, the backscatter regime remains of interest for IPI due to the ability to measure with a single optical access. The use of a single optical access makes IPI much more flexible and opens up the possibility of measurements closer to the actual application rather than laboratory models. In this context, the present work contributes towards a more straightforward and robust applicability of IPI for applications with limited optical access. Complementary to the presented approach, the development of advanced data evaluation methods that are more robust to low SNR appears promising for achieving further progress in IPI as a single optical access application. Here, the primary objectives might be the detection of particles that are less distinguishable from the background, and more robust and accurate algorithms for the fringe frequency extraction."[SDK25]

4 Inverse Problem Approach for the Determination of Particle Location and Size

This chapter addresses RQ3, on how an inverse problem approach can enhance the determination of particle position and size. The proposed methodology integrates two established techniques: DPTV and IPI. Specifically, it combines the 3D position determination capabilities of DPTV with the particle sizing functionality of IPI. The results presented in this chapter build upon the article *Particle Localization with DPTV and Sizing with IPI using a Forward Model and Optimization* by Sax *et al.* [SGK25].

4.1 Combination of Defocusing Particle Tracking and Interferometric Particle Imaging

While promising, the combination of DPTV and IPI introduces several technical challenges that must be addressed to enable accurate volumetric measurements. In IPI, the distance between the particle and the observer directly influences the observed fringe pattern, as encoded in $B_{\text{tot}}(z_P)$ (see Eq. (3.4)). This uncertainty is discussed in Chapter 3, and is estimated in Eq. (3.13). For most optically homogeneous media, $B_{\text{tot}}(z_P)$ varies linearly with the z -position. Consequently, any error in z propagates linearly into the particle diameter estimation, $d_P \sim B_{\text{tot}}$, compare Eq. (3.4). This error propagation into the diameter estimation effectively diminishes the advantage of IPI over other imaging techniques. Due to this sensitivity, IPI is typically restricted to two-dimensional planes where the z -position is known a priori (e.g., via a thin laser sheet). Although recent advances have extended IPI to 3D using APTV [181, 212], the fundamental issue of error propagation from z -position uncertainty remains unresolved. The accuracy of both DPTV and IPI is further compromised by PI overlap. Overlapping PIs can obscure fringe patterns, making it difficult to extract reliable frequency information. This not only affects sizing but also introduces additional uncertainty in z -position estimation. While optical compression techniques [127] can reduce overlap, they are limited to fixed z -planes and are, therefore, unsuitable for volumetric applications. Additionally, the z -position uncertainty in DPTV is further exacerbated by PI overlap, which further complicates PI size estimation, due to even larger uncertainties in $B_{\text{tot}}(z_P)$.

Volumetric implementations of IPI and DPTV face additional constraints related to the usable defocus range. At large defocus distances, the PI diameter grows, causing the light intensity to spread over a larger area. This leads to a reduction in SNR and increases the likelihood of PI overlap, both of which degrade measurement accuracy of both z_P and d_P , and limit the maximum measurable z -position. Conversely, there is also a minimum defocus distance required to avoid aliasing in the fringe pattern. When particles are too close to the focal plane—or when imaging large particles—the fringe spacing becomes too fine to resolve, making both fringe counting and frequency-based sizing unreliable. This minimum defocus distance is not fixed; it depends on the particle size, as the number of observable fringes scales with the particle diameter (d_P). As a result, the effective defocus range for accurate IPI and DPTV measurements is bounded on both ends, posing another challenge for volumetric applications.

Despite these limitations, alternative imaging techniques, such as digital in-line holography, have demonstrated that inverse problem (IP) approaches can overcome similar constraints [61, 185, 129]. Building on these advances, this chapter investigates how such an IP approach can be applied to the combined DPTV-IPI system to enhance both localization and sizing accuracy in three-dimensional particle fields.

"In this type of approach, a forward model (FM) simulates the PIs according to certain particle parameters and the resulting synthetic image is being compared to the observed image. The match between simulated and observed image is quantified by a mismatch functional, which is minimized in the process of finding the unknown particle parameters. [...] The IP approach to DPTV and IPI addresses the following main issues: First, the lower out-of-plane accuracy of DPTV is addressed. Decoupling of the size and z -position estimation prevent error propagation and allows for IPI to be expanded to 3D volumes without loss of sizing accuracy. The issue of PI overlap common to DPTV and IPI is addressed, as the FM can account for superpositions of interference patterns and PIs. Lastly, the aliasing issue is addressed, as the FM can also account for aliasing, in contrast to conventional evaluation methods in IPI." [SGK25]

4.2 The Forward Model - Direct Problem

"To perform IPI with an IP approach, the first step is to build a FM that is sufficiently accurate to the extent that the model error is significantly less than the measurement error. A spherical particle producing a PI has four parameters of interest described by the vector $\psi_P = (x_P, y_P, z_P, d_P)^T$, which are the three-dimensional coordinates of the position and diameter of the particle. The FM should produce a simulated PI based on the particle parameters ψ_P and additional hyperparameters H unrelated to the particle (i.e., wavelength of light λ , focal length of the imaging lens f_{lens} , parameters describing the optical system M_{tot} , etc.). In the following, the FM consists of two cascaded sub-models: A first model for the light scattering at the particle, and a second subsequent model for the wave propagation from the particle to the camera chip through the optical system." [SGK25]

4.2.1 Scattering Model

"The scattering process at the particle is influenced by the particle diameter d_P and the two hyperparameters θ and m . For a large field of view, the scattering angle can furthermore become a function of the particle position. Assuming that the scattering plane is in the x - z plane, the scattering angle becomes $\theta(x_P, z_P) = \theta_0 + \tan(x_P/(z_{P2F} + z_P))$, which is a function of the particle position with θ_0 being the scattering angle of a particle on the optical axis ($x_P = 0$). For the propagation model, three important details about the GPs must be provided by the scattering model: The position of the GPs in the object plane, the intensity of the light at each GP and the phase difference between the two GPs. The scattering process is best described by the LMT, a solution of the Maxwell equations for a homogeneous sphere [200]. However, as determining the GP position with the LMT is computationally expensive, the GP positions and their phase differences is determined using the GO model. The LMT will then just be used to identify the intensity $I = |S_{GP}|$ of the light at these GPs, with S_{GP} being the complex amplitude of the light wave emitted from the GP." [SDK25] A model for determining the position $w^{(p)}$ of a glare point (GP) of scattering order p is provided by Eq. (2.13), in conjunction with the expression $w^{(p)} = q \cos(\beta_i^{(p)})$. To compute the GP position, the parameters k and q in Eq. (2.13), as well as the angle of the ray to the surface tangent $\beta_i^{(p)}$, must be identified. "A fast algorithm for this calculation has

been proposed by Sax *et al.* [SDK25]. This algorithm is used in the present work to compute the positions of the GPs w_j of each order $p = j$ with respect to the centre of the particle projection, see Fig. 4.1.




Figure 4.1: "Scattering model (left) and its integration into the propagation model (right) visualised. Incident light is scattered at the particle. The scattering orders $p = 0$ and $p = 2$ and the position of the GPs w_0 and w_2 are shown. The distance travelled from the incident plane to the object plane is in the surrounding medium l_{out} or inside the particle l_{coord} . The GPs are then modelled as point emitters in the source plane. The length l_{out} allows both GPs to be modelled in the same plane." [SGK25]. Figure from [SGK25].

The second quantity to be provided by the scattering model is the phase difference of the wave between the GPs. For the propagation model, the GPs of a particle are assumed to be in the same x - y plane. Therefore, to obtain the phase difference, the different paths of the light rays from the incoming plane to the object plane can be used, taking into account the refractive index of the respective medium. Knowing the angles β_i and β_t , the travelled distance of the light can be calculated again from optical considerations for $p = 0$ via

$$l_0 = d_P \frac{\sqrt{1 - w_0^2}}{n_2} = 2l_{\text{out}}^{(0)} \quad (4.1)$$

and the refracted orders $p > 0$ via" [SGK25]

$$l_p = d_P \left(\frac{\sqrt{1 - w_p^2}}{n_2} + p \frac{\sqrt{2(1 - \cos(2\beta_t))}}{2n_1} \right). \quad (4.2)$$

In addition, the 180 degree phase change of a reflection on a medium of higher optical density must be considered. "For a reflection on the outside of the sphere, as in the case of $p = 0$, the phase kick-back function $\Delta\varphi_{\text{external}}$ is

$$\Delta\varphi_{\text{external}} = \begin{cases} \pi, & m > 1 \\ 0, & \text{otherwise.} \end{cases} \quad (4.3)$$

In case of an internal reflection as for scattering orders $p > 0$ the function is

$$\Delta\varphi_{\text{internal}} = \begin{cases} \pi, & m < 1 \\ 0, & \text{otherwise.} \end{cases} \quad (4.4)$$

The total phase difference is then computed from

$$\Delta\varphi_0 = 2\pi \frac{l_0}{\lambda} + \Delta\varphi_{\text{external}} \quad \text{for } p = 0, \quad (4.5)$$

$$\Delta\varphi_1 = 2\pi \frac{l_1}{\lambda} \quad \text{for } p = 1 \text{ and} \quad (4.6)$$

$$\Delta\varphi_p = 2\pi \frac{l_p}{\lambda} + (p-1)\Delta\varphi_{\text{internal}} \quad \text{for } p > 1. \quad (4.7)$$

This model provides the phase difference between the GPs. However, it should be noted that this GO model neglects the effect of surface waves and the resulting effect on the phase of the GPs. Therefore, the GO-based phase model is a slight oversimplification and the FM will accordingly have some minor deviations from reality. The phase of the wave affects the position of the stripes of the interference pattern in the PI. However, the frequency of the interference pattern, in contrast, is not affected by the phase. Therefore, in the following, only this frequency is considered as the main feature to determine the particle diameter d_P . For a more accurate phase model, the position of the stripes could also be used to fine tune d_P .

Finally, the intensity of each GP depends on the size of the particle, the scattering angle and the refractive index. The GP intensity must be known, because it influences the magnitude of the constructive interference maxima and the destructive interference minima in the PI. The intensity of each GP S_{GP} is calculated from the Debye series expansion (DSE) of the LMT [40, 68]. This can be done efficiently by expressing the derivatives of the Bessel function through recurrence formulae. To compute the GP intensity a custom MATLAB implementation is used for the computation of the Mie coefficients that follows the stable downward recurrence method of Shen *et al.* [182] in combination with Lentz's continued fraction method [117], and the Mie angular functions are computed using the functions proposed by Maetzler [133]. The implementation has been validated against MiePlot [107] on bubbles ($m = 1/1.333$) and droplets ($m = 1.333$) in the range of $d_P = [10, 300] \mu\text{m}$ and $\lambda = 532 \text{ nm}$. [SGK25]

4.2.2 Wave Propagation Model

"The wave propagation model uses the position (x_{GP}, y_{GP}, z_P) , intensity S_{GP} and phase φ of each GP to model the propagation of light waves through the optical system to the camera chip. Since interference and diffraction must be taken into account, a wave model is used. The model chosen for wave propagation is the Huygens-Fresnel integral, which is an approximation of the Rayleigh-Sommerfeld diffraction theory assuming that the source (object) plane and image plane are much further apart than their respective sizes. The Huygens-Fresnel integral calculates the field $G_2(x_2, y_2, z_2)$ in a plane (x, y) (here the camera chip) at a distance z_2 from the source plane (i.e. the position of the particle) based on the field $G_0(x, y, z = 0)$ in the source plane at $z = 0$, see G_2 and G_0 in Fig. 4.2. The propagation model calculates the waves based on the position of the source plane z_P , i.e. the positions of the glare points. Therefore, the particle" [SGK25] radius "is added to z_P to obtain the true position of the particles centre $(z_P + d_P/2)$, compare the right side of Fig. 4.1. While the Huygens-Fresnel integral only describes the propagation of spherical waves in free space, the introduction of a RTM allows the integral to be generalised to arbitrary optical systems [36]. [SGK25] Optical systems that differ in the x - z plane and the y - z plane can be described by the generalized Huygens-Fresnel integral [36, 15], as given in Eq. (2.51), where $z_{\text{eff}} = \sum_{\ell} z_{\ell}/n_{\ell}$ denotes the effective distance from the source plane to the image plane. The individual distances z_{ℓ} are divided by the refractive index n_{ℓ} of their respective medium to account for the differing optical properties along each segment of the path. "The in-plane coordinates are

(x, y) in the source plane (referring to x_P and y_P) and (x_2, y_2) in the image plane (referring to x_{PI} and y_{PI}). The components $A_{\text{tot}}^x, B_{\text{tot}}^x, D_{\text{tot}}^x$ describe the optical system in the x - z plane and the components $A_{\text{tot}}^y, B_{\text{tot}}^y, D_{\text{tot}}^y$ describe the optical system in the y - z -plane. Eq. (2.51) allows the field G_2 to be obtained even when optical elements such as lenses and glass walls are placed between the source and the image plane. For most optical systems, however, a limiting aperture must be taken into account. A solution to this problem was presented by Wen *et al.* [213, 214], where the problem is solved in two steps: Propagation from the source plane to the principal plane (location of the aperture) and propagation from the principal plane to the image plane, see the notations $G_0, G_1, G_1 T, G_2$ in Fig. 4.2. Starting with the field in the source plane $G_0(x, y, z = 0)$, the field at the aperture $G_1(\xi, \eta, z = z_{\text{aperture}})$ is calculated according to Eq. (2.51). The field G_1 is then cropped by the aperture, which is modelled by multiplying G_1 with a transmission function $T(\xi, \eta)$ in the aperture plane. This transmission function is close to one within the aperture radius and zero elsewhere. A model for such an aperture function as a superposition of Gaussian functions was proposed by Wen *et al.* [213, 214] as

$$T(\xi, \eta) = \sum_{\ell=1}^N E_{\ell} \exp\left(-\frac{Q_{\ell}}{(D_a/2)^2}(\xi^2 + \eta^2)\right) \quad (4.8)$$

with the complex coefficients E_{ℓ} and Q_{ℓ} which can be found in [213, 214] and the aperture diameter D_a . The field at the image plane $G_2(x_2, y_2, z = z_{\text{eff}})$ is then calculated again with Eq. (2.51) based on the cropped field $G_1(\xi, \eta, z = z_{\text{aperture}})T(\xi, \eta)$. For this, the RTM $M_{\text{tot}}^{(x,y)}$ is split into two subsystems describing the optical system from the source plane to the principal plane $M_1^{(x,y)}$ and from the principal plane to the image plane $M_2^{(x,y)}$.




Figure 4.2: "Propagation model and the incorporation of the scattering model visualised. Two particles (double emitters) are shown in a simple optical system consisting of a single lens with an aperture. The GPs of each particle are located in their respective source plane, which has a distance of z_P from the focal plane of the optical system. The field in the source plane is G_0 . The lens and aperture are located in the principal plane. The field G_1 in the principal plane is clipped by the aperture function T . The cropped field $G_1 T$ then propagates to the image plane to form the field G_2 which is the final image seen by the camera $G_{\text{image}} = |G_2|$. Also shown are the coordinates of a particle in the source plane (x_P, y_P) and its counterpart (x_{PI}, y_{PI}) of the particle image in the image plane." [SGK25]. Modified figure from [SGK25].

The majority of the light on the particle surface (seen from the scattering angle θ) is caused by the GPs. Following Shen *et al.* [179] a GP can be approximated as a point source of light at a known position, intensity and phase (i.e., the output of the scattering model). The source field for a single GP is then $G_0(x, y, z = 0) = S_{GP} \delta(x - x_{GP}, y - y_{GP})$ with the complex amplitude of the emitted

wave S_{GP} , which includes the intensity and phase of the GP, and the Dirac function δ . The peak of the Dirac function is placed at the coordinates

$$\begin{pmatrix} x_{\text{GP}} \\ y_{\text{GP}} \end{pmatrix} = \begin{pmatrix} x_{\text{P}} + w_x d_{\text{P}} \\ y_{\text{P}} + w_y d_{\text{P}} \end{pmatrix} \quad (4.9)$$

based on the relative position of the GP on the particle surface $w_{x,y}$ provided by the scattering model, see the particle close up in Fig. 4.2. For most IPI applications, a scattering plane in either the x - z or x - y plane is chosen so that the $w_{x,y}$ component outside the plane is zero. The Dirac function greatly simplifies the solution of the integral, and an analytical solution for G_2 can be given directly from the source field G_0 for a single GP [179]:

$$G_2(x_2, y_2, z_{\text{eff}}) = \frac{\exp(i\frac{2\pi}{\lambda} z_{\text{eff}}) \exp(i\frac{\pi}{\lambda} [\frac{D_2^x x_2^2}{B_2^x} + \frac{D_2^y y_2^2}{B_2^y}])}{(i\lambda)^2 \sqrt{B_1^x B_1^y B_2^x B_2^y}} \sum_{\ell=1}^N \frac{E_{\ell} \pi}{\sqrt{\gamma_{x,\ell} \gamma_{y,\ell}}} S_{\text{GP}} \exp(\beta_{\ell} + i\varphi) \quad (4.10)$$

with the following abbreviations

$$\beta_{\ell} := i\frac{\pi}{\lambda} \frac{A_1^x}{B_1^x} x_{\text{GP}}^2 + i\frac{\pi}{\lambda} \frac{A_1^y}{B_1^y} y_{\text{GP}}^2 - \frac{\Theta_x^2}{4\gamma_{x,\ell}} - \frac{\Theta_y^2}{4\gamma_{y,\ell}} \quad (4.11)$$

$$\gamma_{x,\ell} := \frac{4Q_{\ell}}{D_a^2} - i\frac{\pi}{\lambda} \left(\frac{D_1^x}{B_1^x} + \frac{A_2^x}{B_2^x} \right) \quad (4.12)$$

$$\gamma_{y,\ell} := \frac{4Q_{\ell}}{D_a^2} - i\frac{\pi}{\lambda} \left(\frac{D_1^y}{B_1^y} + \frac{A_2^y}{B_2^y} \right) \quad (4.13)$$

$$\Theta_x := \frac{2\pi}{\lambda} \left(\frac{x_{\text{GP}}}{B_1^x} + \frac{x_2}{B_2^x} \right) \quad (4.14)$$

$$\Theta_y := \frac{2\pi}{\lambda} \left(\frac{y_{\text{GP}}}{B_1^y} + \frac{y_2}{B_2^y} \right) \quad (4.15)$$

Eq. (4.10) gives the field in the image plane for a single GP. Particles such as opaque tracer particles in DPTV can be modelled as single emitters and Eq. (4.10) adequately models such a single emitter particle. For particles such as bubbles and droplets in IPI, the particle is modelled as a double (or triple) emitter from its GPs $(x_{j,\text{GP}}, y_{j,\text{GP}})$ with $j = 1, 2, \dots, N_{\text{GP}}$, where N_{GP} is the number of GPs of a particle (usually $N_{\text{GP}} = 2$). The final field of a particle image $G_{2,P}$ is then modelled as a superposition of the complex fields (GPIs) produced by its individual GPs $G_{2,\text{GP},j}$

$$G_{2,P}(x_2, y_2, z = z_{\text{eff}}) = \sum_{j=1}^{N_{\text{GP}}} G_{2,\text{GP},j}(x_2, y_2, z = z_{\text{eff}}); \quad (4.16)$$

see the upper part of Fig. 4.3. To obtain the full image $G_{2,\text{image}}$ of multiple particles, the complex fields $G_{2,P,\ell}$ of individual particles ℓ are superimposed and the absolute value of the result gives

$$G_{\text{image}}(x_2, y_2) = \left| \sum_{\ell=1}^{N_{\text{P}}} G_{2,P,\ell}(x_2, y_2, z = z_{\ell,\text{eff}}) \right| \quad (4.17)$$

with N_{P} number of particles, compare the lower part of Fig. 4.3.




Figure 4.3: "The superposition principle for the construction of a PI from the individual GPIs is visualised (top). The construction of a full camera image with multiple PIs is shown in the lower half (bottom)."[SGK25]. Figure from [SGK25]

If IPI is to be performed in a volume, each particle ℓ might be located at a different z -position. Therefore, does this model contain a separate source plane for each particle ℓ , but only one common image plane (i.e., the location of the camera chip). Accordingly, also the RTM $M_{1,\ell}$ (i.e. source plane to principal plane) might be different for each particle ℓ as it is a function of the z -position of the particle $z_{P,\ell}$.

The RTM M_2 (describing the propagation from the principal plane to the image plane) is the same for all particles."[SGK25]

4.2.3 Validation of the Forward Model

"The sub-models used for the scattering model have already been validated in the literature [201, 179] and the GO model is commonly used to model the scattering process in IPI [98, 90]. Therefore, only the full FM needs to be validated for the accuracy of the PI diameter as a function of the particle's z -position and for the correct representation of the PI's interference pattern for a given particle size.

To validate the PI diameter d_{PI} of the model, a simple single lens system is modelled, similar to the one shown in Fig. 2.4. The system consists of a single lens of focal length $f_{lens} = 105$ mm and finite aperture (aperture number 4 corresponding to the aperture diameter $D_a = f_{lens}/4$) and an image plane located at $z_{I2F} = 200$ mm. The corresponding focal plane is at $z_{P2F} = 221.1$ mm. The values for f_{lens} , D_a where chosen from the validation experiment in order to have realistic parameters. A single emitter particle is simulated at different distances from the focal plane. The resulting gray value intensities across the PI are plotted in Fig. 4.4. The FM shows a good agreement of the PI diameter with the Olsen-Adrian model [146] and with the GO model [179]. It can also be seen in Fig. 4.4 that the loss of intensity with increasing defocus is captured by the model. This means that the intensity of the PI can be used as an additional feature for the IP.

Eq. (4.10) has been validated to accurately represent the interference pattern for droplets [179]. Therefore, the FM is validated here on an experiment with bubbles ($m = 1/1.333$). The images used to validate the model are images from the experiment described in [SDK25].




Figure 4.4: "Gray value intensity distribution of the imaged PI field for different distances z_P from the focal plane as computed with the forward model. For comparison, the theoretical PI diameter according to the Olsen-Adrian model [146] (•) and the GO model from Eq. (2.4) (•) are converted to pixels and added to the diagram." [SGK25]. Figure from [SGK25]




Figure 4.5: "Optical system of the experiment used to obtain the experimental PIs of the bubbles. The bubbles are located at a distance z_P from the focal plane. The focal plane is located in a water tank behind a glass wall. The imaging system is a single lens with an aperture in the principal plane. The camera chip is located in the image plane. The refractive indices of each part of the optical system are given." [SGK25]. Figure from [SGK25]

The optical setup is shown in Fig. 4.5. The bubbles are located at a distance z_P from the focal plane, in a water tank with $n_2 = n_w = 1.333$. The focal plane is located in the water with a distance of $z_w = 208.3$ mm from the glass wall (thickness $z_g = 4$ mm). The glass wall is located at a distance of $z_a = 117.1$ mm from the principal plane. At the principal plane an imaging lens with a limiting aperture is located (Nikon Micro-Nikkor $f_{\text{lens}} = 105$ mm, $D_a = f_{\text{lens}}/4$). Finally the image plane (camera chip) has a distance of $z_{\text{im}} = 211.8$ mm from the principal plane. The camera used in the experiment was a PCO Pixelfly (pixel size $6.45 \mu\text{m}$) and the bubbles were illuminated with a Quantel Evergreen Nd:YAG laser ($\lambda = 532$ nm) at a scattering angle of $\theta = 99^\circ$. The scattering angle of 99° , was chosen as the measurement uncertainty is the lowest at this scattering angle for the given particle sizes [SDK25]. Further information on the experiment can be found in [SDK25].

The imaging lens is modelled as a thin lens in the principal plane. The effective optical path length of the system is $z_{\text{eff}} = (z_{\text{im}} + z_{\text{a}})/n_{\text{a}} + z_{\text{g}}/n_{\text{g}} + (z_{\text{w}} + z_{\text{p}})/n_{\text{w}}$, where z are distances and n refractive indices. The corresponding RTM from lens to camera chip is

$$M_2^{x,y} = \begin{bmatrix} 1 & z_{\text{im}} \\ 0 & 1 \end{bmatrix} \quad (4.18)$$

and the RTM from particle to lens is

$$M_1^{x,y} = M_{\text{lens}} M_{\text{p,a}} M_{\text{I,g} \rightarrow \text{a}} M_{\text{p,g}} M_{\text{I,w} \rightarrow \text{g}} M_{\text{p,w}} M_{\text{p,P}(z_{\text{p}})}, \quad (4.19)$$

where $M_{\text{p,j}}$ represents propagation through free space, $M_{\text{I},n_1 \rightarrow n_2}$ represents a flat interface from medium n_1 to medium n_2 and M_{lens} describes a thin lens. The indices a, w, g represent air, water and glass respectively.

In the validation experiment, the z -position of the bubbles was measured by fitting the PI diameter with a circular Hough transform. The determination of the PI diameter by this method is at least pixel accurate [SDLK23]. For the given optical system an error of $\Delta d_{\text{PI}} = 1 \text{ px}$ results in an uncertainty of $\Delta z_{\text{P}} = 81.4 \mu\text{m}$ according to Eqs. (2.4) and (2.5). A widely spread method to interpret the interference pattern in IPI is the extraction of the fringe frequency by applying a Fourier transformation to the fringe pattern and fitting the frequency peak in frequency space [16, 42, 51]. This method is, therefore, used as the benchmark in the present work. To determine the size of the bubble, the PI was averaged along the stripes of the interference pattern and the resulting signal, as plotted in Fig. 4.6, was processed by an FFT. The signal was zero-padded by a factor of 10. Analogous to [SDK25], the PI was split up into four different segments perpendicular to the stripes (upper-kernel, middle-kernel, lower-kernel and full-kernel). The four image kernels were then processed individually to obtain the experimental variance of the FFT processing. The particle size was determined using Eq. (3.4).




Figure 4.6: "Experimental (left) and simulated (right) particle images for three different bubble sizes (top). In the bottom row the signal from averaging the PI along the stripes of the pattern is shown. Below each graph the used z -position and size of each bubble is displayed. The uncertainty of the size estimation with the FFT approach is also displayed. The size of the particle was only varied within a tenth of the uncertainty of the FFT approach to obtain the simulated image." [SGK25]. Figure from [SGK25]

To obtain the diameter uncertainty of the method, the kernel with the largest frequency was paired with $z_{\text{P}} + \Delta z_{\text{P}}$ and the kernel with the smallest frequency was paired with $z_{\text{P}} - \Delta z_{\text{P}}$. This gives the maximum spread of possible diameters. The determined z -position and diameter were then used in

the FM to simulate the PIs, which are shown in Fig. 4.6. Since the FFT approach has an estimation uncertainty, the bubble diameter for the FM was varied, starting from the mean diameter of the FFT approach, to best fit the experimental PI. The bubble size was varied only within a tenth of the uncertainty interval of the FFT approach. It can be seen that for different bubble sizes a great agreement between the simulated PIs and the real experimental PIs was achieved within the measurement uncertainty of the FFT approach. In the experiment a heptagonal aperture was used which the model fails to capture as the aperture function T assumes a circular aperture. However, the match of PI diameter and the intensity distribution of the fringes match closely as can be seen in Fig. 4.6. Due to the heptagonal shape of the aperture, the outer most stripes of the averaged signal differ slightly. This is, however, an artefact of averaging along the stripes, where the shape of the aperture affects the height of the averaged maxima. Another discrepancy between the FM and the experiment is the location of the Newton-rings (i.e., diffraction at the aperture) within the PI. As the FM model uses the paraxial approximation, the Newton-rings are always centred in the PIs, unlike in reality where the centre of the Newton-rings depends on the in-plane position of the PI. Over all, the FM represents reality sufficiently and can be used for the IP approach."[SGK25]

4.3 Inverse Problem for a Single Particle

"In this section the IP for a single double-emitter particle is considered. The simulation of multiple particles follows in Section 4.4. A single double emitter (bubble or droplet) is characterised by the parameter vector $\psi_P = (x_P, y_P, z_P, d_P)^T$. To optimise this vector, an interior-point algorithm [88] with numerical gradients by means of central differences is applied. This work uses MATLAB's implementation of the *fmincon*, which uses a trust-region interior point algorithm[27], to minimise a mismatch functional $\mathcal{P}(\psi_P)$ that assesses the similarity between the simulated and the real (observed) image G_{obs} . Note that the same methodology can also be applied to a DPTV only approach locating tracer particles. In this case, a particle is modelled as a single emitter and the optimization vector is shortened to $\psi_P^{\text{Tracer}} = (x_P, y_P, z_P)^T$."[SGK25]

4.3.1 Mismatch Functionals for the Optimization

"The IP approach requires suitable mismatch functionals for the optimization. For the position of the particle $(x_P, y_P, z_P)^T$ the *Structure Similarity Index* SSIM [208] was identified to be a robust choice. The particle is modelled as a double emitter, which leads to the optimization vector $\psi_{P,\text{pos}} = (x_P, y_P, z_P, d_P)^T$. An IPI particle is modelled as a double emitter due to the double-circular nature of the PIs and the influence of the particles diameter on the PI shape, see Fig. 4.7. While the global minimum of the SSIM mismatch functional is in the correct position, it appears to be non-smooth and has many local minima. This is due to the stripes of the interference pattern in the PI. To mitigate this problem, the particle image is simulated without stripes. In order to avoid the stripes due to interference, the two glare point images (GPI) $G_{2,\text{GP}j}$ are not superimposed until after the absolute value of the field has been taken, i.e.

$$G_{P,\text{pos}} = |G_{2,\text{GP}1}| + |G_{2,\text{GP}2}|. \quad (4.20)$$

In this way, the complex amplitude is omitted and no interference appears. The resulting PI obtains only the shape of the PI but not the interference pattern, as indicated in Fig. 4.7. The mismatch functional for the particle position therefore goes as

$$\mathcal{P}_{P,\text{pos}}(\psi_{P,\text{pos}}) = -\text{SSIM}(G_{P,\text{pos}}(\psi_{P,\text{pos}}), G_{\text{obs}}) \quad (4.21)$$

with a negative sign as the SSIM approaches a positive maximum ($\mathcal{P}_{P,pos} = 1$) for a perfect match.




Figure 4.7: "Particle images for different particle diameters and distances from the focal plane. Larger particles require greater distances from the focal plane to be sufficiently defocused (i.e., the PI resembles a single circle). PIs without interference, as created with Eq. (4.20), are shown, to visualise the double circle nature of double emitter PIs (cp. also the blue circles). The images without interference are used to calculate $\mathcal{P}_{P,pos}$. The influence of the particle diameter in combination with the z -position on the PI shape becomes obvious, when omitting the interference. Also the effect of aliasing as a function of diameter and z -position can be observed as marked by the red circle. Recall that aliasing can occur in two ways, where either the particle becomes larger with accordingly increasing number of fringes or the PI becomes smaller, such that the remaining PI-intersection for the fringes diminishes." [SGK25]. Figure from [SGK25]

To formulate a mismatch functional for the diameter $\mathcal{P}_{P,dia}$, the fringe frequency should be considered as the main feature. Simply comparing both G_{image} with the observed (i.e. recorded) image G_{obs} for different particle diameters d_P leads to the same problem of a non-smooth objective functional with many local minima, due to the stripes in the PI. Varying the particle diameter changes the fringe frequency, and whenever stripes of the simulated and observed images overlap, there is a local minimum, even if the fringe frequency of the pattern is not the same. In addition, the phase difference between the GPs causes the stripes to move perpendicular to the stripe orientation as the particle diameter increases, creating even more local minima between the fringe frequency multiples. Both of these effects render a simple comparison of the simulated image with the observed image infeasible for a local search and would require an expensive global search. Therefore, the Fourier-transformed images are compared instead. This eliminates the phase-induced shift of the stripes and multiples of the fringe frequency are clearly distinguishable. Since noise in the observed image introduces additional frequencies, the normalised cross-correlation of the two Fourier transformed images is chosen as the mismatch functional. The mismatch functional for the diameter is

$$\mathcal{P}_{P,dia} = -(\mathcal{F}_2\{G_{\text{image}}\} \star \mathcal{F}_2\{G_{\text{obs}}\}), \quad (4.22)$$

where \mathcal{F}_2 denotes a two-dimensional FFT and the \star symbol denotes a normalised cross-correlation. For better accuracy, the simulated and observed images can be zero-padded similar to the FFT approach. However, this increases the size of the processed image in pixel space, so a factor of two zero padding was used in this work to limit the computational cost. This mismatch functional has a global minimum for a matching diameter, but is only monotone and smooth near the

global minimum. Other local minima exist further away from the minimum, making the diameter determination the most challenging part.

Since both mismatch functionals are in the same range, $\mathcal{P}_{P,\text{pos}} \in [-1, 0]$ and $\mathcal{P}_{P,\text{dia}} \in [-1, 0]$ they can be combined to formulate the combined mismatch function $\mathcal{P}_P = \mathcal{P}_{P,\text{pos}} + \mathcal{P}_{P,\text{dia}}$. In this way, the characteristic of the PI shape changing with particle size can be taken into account in addition to the fringe frequency."[SGK25]

4.3.2 Optimization Strategy for a Single Particle

"Optimizing over all four particle parameters by means of a local search likely means a termination in a local minimum rather than the global one, due to the non-convex nature of the optimization function. However, for a reliable parameter estimation, a minimum as close as possible to the global minimum must be found. As discussed in the last section, the mismatch function for particle size is only quasi-smooth near the global minimum. Therefore, initialisation close to the global minimum is required to ensure reliable convergence of the algorithm. Recent advances in PI detection [48, SDLK23, 35, 56] for DPTV approaches allow for at least pixel accurate placement of bounding boxes (BB) around the PI. The placement of a BB provides values on the PI (x_{PI}, y_{PI}, d_{PI}). Initial values for the in-plane position of the particle (x_P, y_P) can be derived directly from the PI position (x_{PI}, y_{PI}) considering the magnification of the system, similar to conventional PTV approaches. An initial guess for the z -position of the particle can be derived from the PI diameter d_{PI} , e.g., via Eqs. (2.4) and (2.5).

The BB values can then be used to initialise the optimization close to the global minimum. The local search is then constrained to an interval based on the uncertainty of the detection method used. In this work, two BB accuracies are considered: ± 1 px and a more conservative version of ± 5 px. The optimization is constrained to the BB uncertainty to ensure an actual refinement of the initial detection. However, a BB around the PI does not provide information about the particle diameter. Therefore, to obtain an initial estimate of the particle diameter, a global search is performed over the specified diameter interval with a resolution of 50 steps, while keeping the position parameters fixed. In this work the interval of possible diameters $d_P = [20, 300] \mu\text{m}$ is used, where the lower end corresponds to the lower limit of particle diameters measurable with for the given setup (i.e., the minimum diameter for which a full wavelength of the stripe pattern is still in the PI [SDK25]). A global search over the expected range of diameters provides a starting value for the diameter.

As the optimization in high dimensions is difficult, the problem is separated into two sub-problems to reduce the dimensionality of the search: first, only the position (x_P, y_P, z_P) is optimized with d_P as a free parameter, so that the shape of the simulated PI is closer to the observed PI, compare the double circle nature of the PI shown in Fig. 4.7. In this step, only $\mathcal{P}_{P,\text{pos}}$ is used, so that changes in d_P only affect the double circular shape of the PI and stripes are omitted. This bears the advantage of smoothing the mismatch functional, consequently increasing the chances to find a better local minimum.

In the second step, the particle diameter based on the global search is used to initialise d_P as well, using $\mathcal{P}_{P,\text{dia}}$. The optimization then runs over (z_P, d_P) with $\mathcal{P}_{P,\text{pos}} + \mathcal{P}_{P,\text{dia}}$ to refine the z -position and diameter by changing the size, shape and interference pattern of the PI. The in-plane position is kept fixed as it only affects the position of the PI, not the shape and pattern. The particle parameters (z_P, d_P) are optimized together as the double circular shape and interference pattern of the PI depend on both the z -position (size of the circular GPIs and fringe spacing) and on the diameter (distance of the circular GPIs and the fringe spacing). The combined optimization

further reduces the dependence of the diameter determination on the out-of-plane determination, as the two effects that z_P and d_P have on the PIs are not easily separable. While in standard IPI approach, an error in the z -position directly propagates into the diameter estimation without the chance of correction, the combined optimization allows to correct z_P errors while estimating d_P , thus avoiding this issue.

An overview over the full method – from the recording of the raw images to the reconstructed particle field – is shown in Fig. 4.8. First, images are recorded and processed, e.g. by a convolutional neural network (CNN) to obtain BBs of the PIs, compare the work of [48, SDLK23, 56]. The BBs are then used to initialise the parameter vector. The parameter vector ψ_P is then fed into the FM and the resulting image into the mismatch functional. The parameter vector is then optimized to minimize the mismatch functional, e.g. with a trust-region interior-point [27] algorithm. Once the optimization terminates, ψ_P can be used to obtain the particle field. The process works analogue for single or multiple particles."[SGK25]




Figure 4.8: "Scheme of the IP approach in the context of a measurement. First images are recorded and a detection algorithm (e.g. a CNN) is used to generate BBs. The BBs are used to initialise ψ_P and the IP. In the IP ψ_P is fed into the FM to produce a simulated image. The simulated image is then compared to the observed image by the mismatch functional, which quantifies the mismatch by a single value. This value is then minimized (e.g. by a trust-region interior-point algorithm) until a convergence criterion is reached. The particle field is then reconstructed from the final iteration of ψ_P . The IP is analogue for a single or for multiple particles."[SGK25]. Modified figure from [SGK25]

4.3.3 Test on a Single Synthetic Particle Image

"To test the IP approach first synthetic PIs, as generated by the FM, are considered (added Gaussian noise - mean 0.1 and variance 0.1), as the ground truth particle parameters $\hat{\psi}_P$ are known without uncertainty for synthetic PIs. This also allows for particle parameters to be deliberately modified. For all synthetic tests, the optical system to mimic the bubble sizing experiment is used for the FM.

First, only one particle parameter is optimized, while the other three are kept "[SGK25] fixed "(at ground truth) to investigate the potential limit of the IP approach. To investigate the diameter accuracy, the diameter is initiated with the global search and then fitted with the local search. An

FFT algorithm, as used to size the particles in Fig. 4.6 is also applied on the images to act as a bench mark (also with ground truth z -position). Since the IP approach uses the same FM as used for the generation of the (synthetic) observed image, the IP approach has a natural advantage over the FFT approach on synthetic images, that will not appear on images from real experiments. However, the effect of different PIs on the IP and FFT approaches can be assessed independent from this effect. The results are shown in Fig. 4.9a.

Since the particle diameter and its z -position are connected, two cases are tested: First, the particle diameter is kept constant ($d_P = 100 \mu\text{m}$) and the PIs are generated at different z -positions. This way the effect of insufficient defocusing can be investigated (i.e., when the GPIs are not sufficiently overlapping, see Fig. 4.7 thus leading to the aforementioned aliasing effect). The aliasing effect in the interference pattern appears when the distance between the stripes becomes too small to be resolved by the pixel size (i.e., the Nyquist frequency is reached). As indicated in Fig. 4.7, aliasing can occur through two different ways: (i) the particle becomes too large and too many fringes are squeezed into the PI or (ii) the distance to the focal plane decreases and the PI shrinks so that the same amount of fringes becomes squeezed together in a smaller area.

Both the FFT and the IP approach achieve sub-micrometer accuracy. It can be seen that the IP approach determines the diameter at least to the same level of accuracy as the FFT approach. The main difference between the FFT and the IP approach, is that the IP approach fits the whole spectrum of frequencies in Fourier space, whereas the FFT approach only fits one dominant peak. This renders the IP approach potentially more robust to overlaps and other imperfections of the PI. While the optimization on synthetic images reaches errors as low as sub-nanometer levels, this degree of accuracy is unrealistic as the FM is unlikely to represent reality to this degree. Errors smaller than a tenth of a micrometer can be accredited to the IP approach using the same FM model as the observed images.

A major advantage of the IP approach over the FFT becoming visible in Fig. 4.9a (red marking), which is the robustness against aliasing in the interference pattern. It can be seen that the FFT approach fails to function properly when the particle is too close to the focal plane and aliasing occurs, as seen by the sudden increase of error for small z -positions. This sharp increase does not




Figure 4.9: "IP approach and FFT for the diameter measurement with known particle position (x_P, y_P, z_P). Both approaches were given the ground truth particle position and only the particle diameter was determined. For the simulated observed image, in (a) the particle z -position was varied with a fixed diameter, and in (b) the particle diameter at fixed z -position was varied. The red dashed line (---) represents the micro-meter accuracy for the present optical system. The red dot (•) shows where the two test series intersect in parameter space." [SGK25]. Figure from [SGK25]

occur for the IP approach, since the FM can correctly account for the aliasing (due to the known pixel size in the FM and the according discretization of the field G_2). As the FM accounts for aliasing, the IP can deal with this effect accordingly. This has an important implication: Since aliasing occurs for insufficiently defocused particles (i.e., particles too close to the focal plane for their size), this allows the IP approach for the measurement of particles closer to the focal plane.

Next, the effect of the particle size on the diameter estimation was tested, see Fig. 4.9b. In this test series the z -position of the particle was kept constant and the diameter was varied. The z -position $z_P = 7$ mm was chosen to ensure sufficient defocusing for all particle diameters. It can be seen that the FFT approach performs better for larger particles (i.e., more fringes in the PI), while the particle size seems to have no visible effect on the IP approach.

The same two tests (with the same PIs) were conducted with the in-plane position (x_P, y_P) and the diameter d_P kept fixed at known ground truth values and only the z -position was optimized. For this test a BB placement uncertainty of $x_{PI}, y_{PI}, z_{PI} = \pm 1$ px was assumed. Most machine learning methods used for particle detection in DPTV achieve pixel accurate results [48, SDLK23, 56, 135]. However, the higher uncertainty was chosen for conservative testing and to account for any placement errors of detection algorithms. The results are shown in Fig. 4.10. Overall, it can be seen that the IP approach achieves sub-pixel accuracy for the z -position in the case of good initialization.




Figure 4.10: "Optimization of the particle z -position with the ground truth of the other particle parameters (x_P, y_P, d_P) for a varying z -position and fixed diameter (a) and for a varying diameter with fixed z -position (b). The black dashed line (---) represents the pixel accuracy. The red dot (•) shows where the two test series intersect in parameter space." [SGK25]. Figure from [SGK25]

The diameter of the particle, and therefore, the shape of the PI, compare Fig. 4.7, does not seem to have an observable influence on the z -position uncertainty.

Next the IP approach is considered for a search over the full optimization vector ψ_P , without any ground truth information. The position initialized from the BB placement with a ± 1 px and a ± 5 px error in a random direction to simulate the BB placement uncertainty of a detection approach. For this test, 200 PIs were generated with random particle parameters ψ_P in $x_P, y_P \in [-1, 1]$ mm, $z_P \in [1, 8]$ mm and $d_P \in [20, 300]$ μ m. The distribution of the resulting estimation errors are shown in Fig. 4.11 and the median errors are shown in Tab. 4.1. The in-plane accuracy of the IP approach is sub-pixel accurate and poses a clear improvement from the pixel accurate BB placement. For single PIs the accuracy of the x and y positions differs, which is a result of the directionality of the fringes in the PI. The directionality causes the objective functional to

have a different topology in the x and y dimensions. This effect was, however, not observed when optimizing over multiple particles, see Section 4.4.2 and it seems to become irrelevant in case of higher dimensional spaces. The in-plane accuracy of the IP approach is comparable to current bench marks in DPTV such as the cross correlation approach [166] and the circular Hough transform [8]. A poor initialization has only little effect on the in-plane position estimation, due to the smoothness of the mismatch functional for a single PI in x and y direction. For the z -position accuracy, a significant improvement over pixel accuracy is achieved, with the error being in the same-order of magnitude as the in-plane error. The effect is less pronounced for a poor initialization but still significantly better than pixel accuracy. However, this effect might not hold for the application on real images due to differences in the aperture shape between simulation and real experiment.

Method	median error in μm			median Error in pixel	
	(x_P, y_P) [μm]	z_P [μm]	d_P [μm]	(x_P, y_P) [px]	z_P [px]
FFT + BB (1 px)	(6.45, 6.45)	81.4	47.2	(1,1)	1
IP (1 px)	(2.43, 0.38)	3.13	0.16	(0.38, 0.06)	0.04
IP (5 px)	(3.65, 0.47)	42.4	0.85	(0.57, 0.07)	0.52
improvement factor IP (1 px)	(2.6, 17.0)	26.0	295	-	-
improvement factor IP (5 px)	(1.8, 13.7)	1.9	55.5	-	-

Table 4.1: "Median Error for 200 synthetic PIs with random particle parameters ψ_P . The particles have diameters in the range of $d_P = [20, 300] \mu\text{m}$ and a random distance from the focal plane in the range of $z_P = [1, 8] \text{ mm}$. For the FFT approach a random error within an uncertainty of $\Delta x_{PI}, \Delta y_{PI}, \Delta d_{PI} = 1 \text{ px}$ was added to the ground truth parameters ψ_P . The FFT was then performed on the initialised value. For the IP approach a random placement error for the BB $\Delta x_{PI}, \Delta y_{PI}, \Delta d_{PI} = 1 \text{ px}$ and 5 px was used to initialise the starting values of the optimization. The factor of improvement for the IP based on the chosen FFT+BB (with 1 px error) approach is given in the last two lines of the table. The improvement factor of the IP with 5 px initialization error relates to an improvement over the FFT+BB with 1 px error." [SGK25]. Table from [SGK25]




Figure 4.11: "Probability density function (PDF) of the errors of the IP approach (1 px initialisation error) for 200 single particle images with randomly distributed positions and sizes. The particles have diameters in the range of $d_P = [20, 300] \mu\text{m}$ and a random distance from the focal plane in the range of $z_P = [1, 8] \text{ mm}$. The in-plane position was randomly generated for the PI to be still fully within the image. The pixel accuracy is shown as the black dashed line, and the micro-meter accuracy (for the present optical system) as a red dashed line." [SGK25]. Modified figure from [SGK25]

The biggest improvement of the IP approach is in the estimation of particle diameter d_P . As the FFT approach is reliant on a known z -position, even small errors in the z position cause large deviations in the diameter estimation as the error propagates linearly onto the diameter estimation, compare Eq. (3.4). This is not the case for the IP approach and even for a poor initialization ($\pm 5 \text{ px}$) sub-micrometer accuracy is achieved. The correlation between the z -position and diameter errors is 0.06. Even for a less accurate z position ($42.4 \mu\text{m}$), as is the case for the poor initialisation, remains the diameter estimation largely unaffected ($0.86 \mu\text{m}$). The uncorrelated

diameter estimation poses an improvement of the diameter estimation of more than two orders of magnitude for single PIs."[SGK25]

4.3.4 Transferability to Real Particle Images and Robustness of IP Approach

"For synthetic PIs the IP approach shows some improvement of accuracy over the FFT approach. However, the IP approach takes advantage from using the same FM as the synthetic observed images. Therefore, the transferability of these results to real images needs to be investigated."[SGK25]

"To test this the same PIs used for the validation of the FM, see Fig. 4.6, are used. The IP approach is applied to the real images, using an initialization from BBs with a placement accuracy of $\Delta x_{PI}, \Delta y_{PI}, \Delta d_{PI} = \pm 1$ px.

The result of the optimization can be seen in Fig. 4.12. For all three particles the IP approach returns a result within the measurement uncertainty of the FFT + BB (1 px uncertainty) approach. The amount of improvement of estimation accuracy, can however not be quantified, since the ground truth bubble position and diameter is subject of the uncertainty of the conventional measurement method (FFT + BB) and therefore not known with higher accuracy. With the IP approach, a distinct minimum for ψ_P is found within the uncertainty of the FFT + BB approach. This means that the IP approach is at least as accurate on real data as the FFT + BB approach. The higher accuracy on synthetic data (see Sec. 4.3.3) and the distinction of the minimum within the uncertainty of the FFT + BB approach suggest, that the IP approach is considerably more accurate than the FFT + BB approach. Due to the unknown ground truth, however, this improvement is difficult to quantify. The IP approach also achieves sub-micrometer accuracy (for the given optical system) for the diameter even in the case of an unknown z -position, and sub-pixel accuracy of the z -position despite the mismatch of the heptagonal aperture with the assumed circular one in the FM. It is likely that a heptagonal aperture function would further increase the z -position accuracy of the IP approach, but would render the approach less flexible to optical systems with other apertures.

A common issue in IPI is low signal-to-noise ratio (SNR) as occurs for large defocus distances, compare Fig. 4.4, small particles and measurement at certain scattering angles [SDK25]. To test the IP approach for the robustness against low SNR, the same images as used in Figs. 4.6 and 4.12 are used, with white Gaussian noise added to decrease the SNR post-measurement. The original images all had an SNR larger than 30 and noise was systematically added until the SNR was lowered to 2.6, compare Fig. 4.13. The SNR was defined as the average intensity of the peaks in the fringe pattern (i.e. stripes) compared to the variance of the background. Variance is commonly used to quantify noise power independently of shifts in background intensity, particularly under the assumption of Gaussian-distributed noise, which is a standard model for image noise. Since the ground truth of the bubbles in the real experimental images is only known with the uncertainty of the FFT + BB approach, the error with increasing noise is quantified by the consistency of the IP approach returning the same result for ψ_P despite the added presence of noise. The accuracy reference ($\hat{x}_P, \hat{y}_P, \hat{z}_P, \hat{d}_P$) is therefore the result returned by the IP approach for the original image (without added noise, SNR>30). The error stemming from added noise is then defined as the difference of the returned values of the IP on images with noise (x_P, y_P, z_P, d_P) from the values of the IP on the original images.




Figure 4.12: "Real PIs from the bubble experiment [SDK25] (left), the PI from the IP result (middle) and the image overlay (right) (observed - blue channel; IP - red channel). Given are the ground truth values \hat{z}_P, \hat{d}_P (i.e., measurement with conventional approach) with the measurement uncertainty (left), the values obtained from the IP approach z_P, d_P (middle) and the discrepancy of the two methods (right). The same images as in Fig. 4.6 for the validation of the FM were used. Similar to Fig. 4.6, the blue line on the right represents the experiment and the orange line, the reconstruction by the IP." [SGK25]. Figure from [SGK25]

The results are shown in Fig. 4.13, which outlines that all dimensions of optimization behave very similar in the presence of noise. For higher SNR levels almost no influence of noise is observed for the 3D localization, while for extremely low SNR levels of close to 2.6, the returned result starts to vary more significantly. A closer look at the topology of the mismatch-functional reveals that the smoothness of the optimization landscape is barely affected by the noise, but the global minimum experiences a small shift due to the noise. This reveals that the influence of noise on the IP is an issue affecting the mismatch functional and is, therefore, not an FM or solver-related issue. The increased variance of ψ_P in the presence of noise can therefore, be explained by noise introduced shifts of the global minimum. However, despite very low SNR, the returned result stays sub-pixel accurate for the position and sub-micrometer accurate for the particle diameter determination." [SGK25]




Figure 4.13: "Consistency of the IP approach for images with increased noise levels. Images from the bubble experiment were used and Gaussian noise was added to decrease the SNR (top). The change of the optimization result of the IP approach for different noise levels is shown for the x -position (left), the z -position (middle) and the diameter (right). The values $(\hat{x}_P, \hat{z}_P, \hat{d}_P)$ refer to the results of the IP approach without added noise, see Fig. 4.12. The error is the deviation from the 'noise free' results. The same images as in Fig. 4.6 for the validation of the FM were used." [SGK25] Figure from [SGK25]

4.4 Inverse Problem for Multiple Particles

"The last section demonstrated the feasibility of the IP approach for single particles and suggests improvements in measurement accuracy and noise resistance compared to current benchmark approaches. Almost all measurement scenarios, however, involve multiple particles. In such multi-particle systems, the PIs must be considered together, as the PI overlap in turn leads to superimposed patterns. These overlap regions of the PIs, affected by the superposition, can become complex as multiple fringe frequencies and orientation of the fringes can occur. To accurately fit a PI onto another PI, the other PI's parameter must be known beforehand. Therefore, the IP approach for multiple particles optimizes all particles in the same image at once. Similar to the single particle approach, machine learning techniques for object detection can again be used to place BBs around every PI to provide starting points for the x, y, z positions of each particle. "[SGK25]

4.4.1 Mismatch Functionals and Optimization

"To fit all particles in the image at once, the particle parameter vector is extended to

$$\psi_I = \underbrace{(x_{P,1}, y_{P,1}, z_{P,1}, d_{P,1})}_{\text{particle 1}}, \underbrace{(x_{P,2}, y_{P,2}, z_{P,2}, d_{P,2})}_{\text{particle 2}}, \dots, \underbrace{(x_{P,N}, y_{P,N}, z_{P,N}, d_{P,N})}_{\text{particle N}}^T \quad (4.23)$$

with the length $4N$ for N particles in the image ($\psi_I = (\psi_{P,1}^T, \psi_{P,2}^T, \dots, \psi_{P,\ell}^T, \dots, \psi_{P,N}^T)^T$). The image G_{image} is then computed from the complex fields $G_{2,\ell}$ of the individual particles ℓ using Eq. (4.17). The number N of particles in the image is assumed to be known and can be retrieved from a particle detection algorithm (i.e., number of BBs in the image), compare the initialisation scheme in Fig. 4.8. While, there is a lot of research (e.g. [56, 35, 135]) published on the position accuracy of neural networks, only few papers discuss the detection rate [48, 159, 223, SDLK23]. However, recent approaches for particle detection algorithms have shown to have sufficiently low miss rates to initialise an IP approach. To deal with the many local minima in this high dimensional space, again the initialization based on the BBs $(x_{PI,\ell}, y_{PI,\ell}, d_{PI,\ell})$ of each PI_ℓ is used to obtain good starting values for the optimization.

The optimization routine remains identical to the single particle optimization: first the positions are optimized and in the second step the z -positions and the diameters are refined. In the first step the full vector ψ_I is fitted in order to identify the individual particle positions and the diameter is again kept as a free parameter in order to adapt the PI shapes. To suppress interference in this step the image is computed from

$$G_{\text{image,pos}} = \sum_{\ell=1}^{N_P} G_{P,\text{pos},\ell} = \sum_{\ell=1}^{N_P} |G_{2,\ell,\text{GP1}}| + |G_{2,\ell,\text{GP2}}| \quad (4.24)$$

for N_P particles. The mismatch functional $\mathcal{P}_{I,\text{pos}}$ is then computed identical to the single particle approach ($\mathcal{P}_{P,\text{pos}}$), compare Eq. (4.21). To compute the mismatch functional for the diameter, the PIs are Fourier-transformed individually. For this, an image snippet of each particle ℓ is used for the mismatch functional. The image snippet is computed from a BB, of which the centre is derived from $x_{P,\ell}$, $y_{P,\ell}$ and the BB diameter is derived from $z_{P,\ell}$ and Eq. (2.4). The mismatch functional $\mathcal{P}_{I,\text{dia}}$ is then computed from the individual Fourier-transformed image snippets and takes the form of

$$\mathcal{P}_{I,\text{dia}} = -\frac{1}{N} \sum_{\ell=1}^N (\mathcal{F}_2\{G_{\text{PI},\ell}\} \star \mathcal{F}_2\{G_{\text{PI,ref},\ell}\}), \quad (4.25)$$

which is an expansion of Eq. (4.22). To initialize the diameters, again a global search is run over the individual PIs using Eq. (4.22). In the second step the local search again only optimizes the z -positions and diameters so that the optimization vector in the second step takes the form of

$$\psi_{I,\text{zd}} = \underbrace{(z_{P,1}, d_{P,1})}_{\text{particle 1}} \underbrace{(z_{P,2}, d_{P,2})}_{\text{particle 2}} \dots \underbrace{(z_{P,N}, d_{P,N})}_{\text{particle N}}^T \quad (4.26)$$

and has the size of $2N$. Since the z -positions and diameters of the particles are optimized together, the objective functional of the second step takes again the form of $\mathcal{P}_I = \mathcal{P}_{I,\text{dia}} + \mathcal{P}_{I,\text{pos}}$, similar to the single particle approach. The approach for multiple particles is essentially identical to the single particle approach, but the generated images contain multiple PIs.

The multi-particle IP approach has the following advantages: The shape of overlapping PIs and the super-positioned fringe pattern can be accurately represented by the FM. The LMT-DSE model in the FM intrinsically models changes in intensity with varying particle size and scattering angle (i.e., x, y position of particle). Also the intensity change with the z -position is accounted for, which is indicated by the intensity loss in Fig. 4.4. The increased intensity and varied interference pattern in overlapped parts between PIs is also accounted for by the FM as the complex fields of the PIs are superimposed. These characteristics let the approach use the relative intensity between PIs as an additional feature, which other approaches do not take advantage of."[SGK25]

4.4.2 Tests on Synthetic Data

"To test the multi-particle IP approach, again synthetic images are generated with the FM to have a ground truth for the analysis. Similar to the single particle tests, noise is added to the observed images and the optimization is initialized from BBs with ± 1 px and ± 5 px placement uncertainty (random derivations for each BB $\Delta x_{\text{PI}}, \Delta y_{\text{PI}}, \Delta d_{\text{PI}} \in [-5, 5]$ or $[-1, 1]$). For comparison, also the FFT + BB (± 1 px uncertainty) approach is applied to the images. The random initialization for the FFT + BB approach and for the IP approach where exactly the same (for the IP approach multiplied by 5 in case of the ± 5 px initialization) to cancel out the influence of 'lucky' placements of BBs. Ten different images containing 20 particles each were generated for testing.

The results of the test are shown in Fig. 4.14. It can be seen that the multi-particle IP approach remains sub-pixel accurate in the in-plane position for a good initialization, but not for a poor initialization (± 5 px), see the improvement factors in Tab. 4.2. This is due to the non-smooth landscape of the optimization function, where the optimization gets stuck in a local minimum, if not initialized properly. In the case of a poor initialization (± 5 px) the in-plane error is still reduced from the BB error of $32.25 \mu\text{m}$ to $10.5 \mu\text{m}$, which poses an improvement by the factor of 3.




Figure 4.14: "Probability density functions for the in plane x and y errors (a), the out-of-plane errors (b) and the diameter estimation error (c). For the diameter error also the FFT approach is plotted as comparison. Similarly to the FFT, the IP approach was initialised with 1 px uncertainty. Ten images with 20 particles each were generated with random particle parameters in $d_P \in [20, 300] \mu\text{m}$ and $z_P \in [1, 8] \text{mm}$. The pixel accuracy is shown as the black dashed line, and the micro-meter accuracy as a red dashed line. The yellow dashed line in (c) marks the threshold beyond which outliers are defined." [SGK25]. Figure from [SGK25]

Method	median Error			number of outliers ($ \hat{d}_P - d_P > 10 \mu\text{m}$)
	$(x_P, y_P) [\mu\text{m}]$	$z_P [\mu\text{m}]$	$d_P [\mu\text{m}]$	
FFT + BB (± 1 px)	(6.45, 6.45)	81.4	34.4	25
IP (± 1 px)	(1.68, 1.77)	3.15	0.72	5
IP (± 5 px)	(10.5, 14.4)	61.8	2.81	10
improvement factor IP				
(initialization ± 1 px)	(3.8, 3.6)	25.8	47.7	5
(initialization ± 5 px)	(0.61, 0.44)	1.3	12.2	2.5

Table 4.2: "Median errors of the FFT + BB (1 px uncertainty) and the IP approach (initialisation with 1 px and 5 px uncertainty). The factor of improvement of the result using the IP approach is given below. The number of outlier is stated, with an outlier, arbitrarily being defined as an error in diameter of more than $10 \mu\text{m}$. The outlier threshold of $10 \mu\text{m}$ was chosen based on the PDFs in Fig. 4.14c, where two ranges of high probability are separated: one in the order of magnitude of $1 \mu\text{m}$ and the other in the order of magnitude of $100 \mu\text{m}$. The improvement factor of the IP with 5 px initialization error relates to an improvement over the FFT+BB with 1 px error." [SGK25]. Figure from [SGK25]

The z -position error remains almost identical for the multi-particle approach (3.13 vs. $3.15 \mu\text{m}$ error for the good initialization and 42.4 vs. $61.8 \mu\text{m}$ for the poor initialization). The landscape of the mismatch-functional is mostly convex in the z dimension for single particles and is also smoother than in the x, y dimensions for multi-particle systems. The IP approach improves the particle parameter compared to its initialization significantly even for a poor initialization. For the diameter estimation also similar results to the single particle approach are achieved. The IP approach achieves significantly better results in the diameter estimation due to the uncorrelated diameter and z position determination. The investigation of the diameter error shows that the error probability density function (pdf) is split into two parts. An arbitrary threshold of a diameter error of the order of magnitude of $10 \mu\text{m}$ and above is defined to mark outliers in the estimation,

based on the separation of the two pdfs, compare Fig. 4.14c. The number of outliers is significantly reduced with the IP approach compared to the FFT+BB approach (25 to 5). "[SGK25]

4.4.3 Particle Image Overlap

"As the IP approach achieves very similar results for multiple particles compared to single particles, the effect of PI overlaps on the result is investigated in more detail. To investigate the influence of PI overlaps in images with multiple particles, the cumulated intersection over area (IoA) of each PI is defined. The IoA describes the amount of overlap between two objects in relation to the area of the considered PI. It therefore shows the fraction of an individual PIs area that is covered by another particle. The IoA is introduced as

$$\text{IoA}_{j,\ell} = \frac{A_{\text{PI},j} \cap A_{\text{PI},\ell}}{A_{\text{PI},j}} \quad (4.27)$$

with $A_{\text{PI},j}$ being the area of the PI. The cumulated IoA for a particle j sums up the IoA of the PI j with every other PI ℓ in the image

$$\text{cumulated IoA}_j = \sum_{\ell} \text{IoA}_{j,\ell} \quad (4.28)$$

and, therefore, provides a measure for the total amount of area in the PI that is covered up by overlap. Fig. 4.15a shows the z -error for each PI in relation to the cumulated IoA. From Fig. 4.15a, it can directly be seen that the z -error is unaffected by the overlap. The same goes for the diameter error as can be seen in Fig. 4.15b. The comparison between the FFT and the IP approach, both for ± 1 px BB placement accuracy, reveals that the IP approach produces significantly less outliers and achieves consistently lower errors in the presence of strong overlaps.




Figure 4.15: "The z -position (a) and diameter (b) errors over the cumulated Intersection over area (IoA) of the PIs. The IoA is the intersection area of a PI divided by its area. The accumulated IoA is the summed IoA of a PI with all other PI overlaps and shows the amount of the PI area that is overlapped. In (a) the dashed black line represents the pixel accuracy. In (b) the thresholds of a $1\mu\text{m}$ diameter error (red line) and for the outlier definition of an error of $10\mu\text{m}$ (yellow line) are visualised. The errors are shown for the IP approach and for the FFT approach both with a 1 px BB accuracy for initialisation. The FFT approach and the IP approach are initialised from a ± 1 px BB placement accuracy." [SGK25]. Figure from [SGK25]

To visualise the performance of the FFT+BB approach compared to the IP approach, Fig. 4.16 shows two of the test images with the respective error category (< 1 , $[1, 10]$ and $> 10\mu\text{m}$ error). It can be seen that the FFT approach (fitting a single peak in frequency space) works well for non-overlapping PIs but creates a significant amount of outliers in the presence of even small PI

overlaps. The FFT approach results in large errors even for small amounts of overlap. This is due to the presence of other significant peaks in the Fourier space due to the superposition of frequencies. Consequently, any of the frequencies might be detected and it is difficult to identify the correct one. For the IP approach on the other side, most particles could be reconstructed despite larger PI overlaps. The number of outliers in the presence of overlaps is reduced from 25 (FFT) to 5 with the IP approach. However, the IP approach still produces some outliers. These outlier PIs were all significantly less bright than the other PIs surrounding them, indicating that the intensity of PIs does have an important role on the reconstruction. This is likely due to brighter PIs influencing the mismatch functional more than dimmer ones. PIs less bright than their surrounding PIs stem either from smaller particles or more defocused particles. This means that these two types of particles are most likely to be wrongly reconstructed.




Figure 4.16: "The estimation results of two of the 10 synthetic test images visualized for the FFT + BB (1 px placement uncertainty) and the IP approach. A colored dot is added on each PI to show the error of the diameter determination: For an error $< 1 \mu\text{m}$ a green dot is placed (\bullet), for an error in $[1, 10] \mu\text{m}$ a blue dot (\bullet) and for outlier particles (error $> 10 \mu\text{m}$) a red dot is placed (\bullet)."[SGK25]. Figure from [SGK25]

Considering the many outliers caused by PI overlaps when using the FFT approach, the IP approach yields a promising alternative for the improvement of median and mean accuracy in diameter estimation. However, the IP approach also returns larger errors on PIs, which are isolated from other PIs without any overlap. This stems from the challenging topology of the mismatch-functional."[SGK25]

4.4.4 Particle Reconstruction from the IP Approach

"Finally, the IP approach is used to reconstruct a full three dimensional particle field, from the (synthetic) observed image. For the visualization of the application, one of the ten test images from




Figure 4.17: "Synthetic observed image and the reconstruction by the IP approach visualized. The reconstructed particles are compared to the ground truth used to generate the observed image. In the present example, one particle was reconstructed with a false diameter but correct position, and is marked by the red circle." [SGK25]. Figure from [SGK25]




Figure 4.18: "Example of the use of the IP approach on real experimental data. The observed image is from the bubble experiment. The observed image is fitted by minimizing the mismatch functional \mathcal{P} . The FM is used to generate an image with the optimal particle parameter set ψ_I . This image can be used as a reference to qualitatively check the result of the IP approach. The particle positions and diameters are then directly obtained from ψ_I ." [SGK25]. Figure from [SGK25]

Section 4.4.2 is chosen. Fig. 4.17 shows the ground truth particle field with the three dimensional positions and diameters, that was used to generate the observed image with the FM. The reconstructed particle field is shown next to the ground truth. It can be seen that all but one particle were reconstructed with sufficient accuracy. The position of this one particle was still reconstructed correctly and only the diameter was estimated wrongly. Finally, Fig. 4.18 shows the application of the IP approach on a real image from the bubble experiment [SDK25], which was also used to validate the FM in Fig. 4.6, with multiple particles. The IP approach finds the particle parameters ψ_I that result in a good agreement of the observed image with the reconstructed image according the mismatch functional \mathcal{P} . For a qualitative check of the reconstruction, the reconstructed image, based on the determined ψ_I , can be generated using the FM. A visual comparison shows great agreement between the experimental image and the reconstructed image. The particle field is then directly obtained from the particle parameter vector ψ_I ." [SGK25]

4.5 Discussion and Conclusions

"There are several challenges in DPTV and IPI that the proposed IP approach aims to address. The first is the out-of-plane accuracy in DPTV, which is usually an order of magnitude lower than the estimation of the in-plane position. PI overlaps furthermore present a limit to the number of particles that can be evaluated, the maximum defocus length and defocus sensitivity, as well as the spatial resolution of fringes in IPI. Moreover, the coupling of the z -position with the diameter estimation present an issue in IPI for the sizing of 3D distributed particles. Lastly, aliasing presents a lower limit on defocusing and upper limit on the particle size in IPI.

The inclusion of a FM in the IPI measurement offers several advantages: It allows an approach to use the relative intensity of PIs as an additional feature to extract information. Effects such as changing intensity due to particle size, z -position and scattering angle (x, y -position) [SDK25] are taken into account by the FM, which would otherwise have to be considered manually. This is particularly important for IPI in side- and backscatter where large intensity variations can occur.

The IP approach showed to have a z -position accuracy in the same order of magnitude as the in-plane position on synthetic data. A similar trend was observed on real data, however the exacted uncertainty on real data is difficult to quantify. The IP approach, therefore poses a promising counteraction to large out-of-plane uncertainties in the 3D position determination with a single camera by means of defocusing.

Aliasing effects of the fringe pattern can also be taken into account. There are two ways aliasing can occur, by either large particles or by particles close to the focal plane, see Fig. 4.7. According to Eq. (3.4) and B_{tot} being linearly depended on z_P , is the Pareto front of occurring aliasing linear. Therefore, an increase in d_P must be compensated by a linear increase in z_P to avoid aliasing. While conventional FFT based approaches are limited by aliasing, the IP approach can operate beyond the point of aliasing (decreased z_P for a given particle size d_P). As a result, the focal plane can be moved closer to the particles in the measurement, resulting in less overlap and higher intensities (i.e., higher SNR). This is a significant advantage over non-FM approaches as the quality of the raw data can be improved by imaging closer to the focal plane. This is particularly important for measurements in the side- and backscatter regions, which generally suffer from low SNR [SDK25]. The approach presented, therefore, addresses one of the most important issues in backscatter IPI. In addition, the IP approach shows great resilience to the presence of noise, to the point where most particle detection methods fail. The ability of the IP approach to deal with low SNR levels (e.g., from particles further away from the focal plane) and aliasing (e.g., from particles closer to the focal plane) allows the measurement volume in the z dimension to be extended in both directions. Consequently, potentially deeper volumes can be measured with the IP approach compared to other methods. Alternatively, this characteristic of the IP approach can also be used to allow for the use of optical systems with higher defocusing sensitivity (larger change in d_{PI} with z_P) to further increase the z -position accuracy.

Comparisons with the FFT+BB approach show that the IP approach is at least as accurate as the current benchmark on real data with sub-micron accuracy in diameter measurement for the optical setup in this work. The IP approach has another significant advantage over the FFT approach in that the z position and the diameter estimation are decoupled. This results in significantly higher diameter accuracy for three dimensionally distributed particles, as the error in the z -position estimation doesn't propagate into the diameter estimation. This feature makes the IP approach more accurate than sequential approaches that first determine the z -position of a particle and then use this estimate to calculate the diameter. The IP approach presented in this work introduces a competitive approach for the simultaneous position and size determination of three-dimensionally distributed particles, alongside currently existing approaches such as IPI+APTV approaches [180, 181, 212] and IP in holography [185, 61]. The presented IP approach can be applied to systems with optical elements such as lenses, apertures and glass walls, which are challenging for holographic approaches. The approach can also be used to size both droplets and bubbles, as the scattering model takes into account the different scattering process for any real-valued relative refractive index.

The IP approach can be extended to measure multiple particles simultaneously. The main problem in multi-particle systems is PI overlap, which causes approaches that rely solely on frequency extraction (either by FFT or counting fringes) to produce erroneous estimations due to the presence

of other frequencies caused by the superposition of PIs. The IP approach, however, is largely unaffected by PI overlap. The approach, therefore, addresses the PI overlap problem in both DPTV and three-dimensional IPI. The IP approach, therefore, relaxes the trade-off between the need for larger PIs to improve fringe resolution (and avoid aliasing) and the need for smaller PIs to reduce PI overlap [127]. Another factor affecting this trade-off is the need for high SNR, which again makes small PIs advantageous. With a high noise resistance and intensity as an additional feature in the FM, can the IP approach also deal with larger defocused PI with lower SNR, and, therefore, partially mitigates this part of the trade-off as well.

While the IP approach offers advances in many of the key challenges in DPTV and IPI, the accuracy of the approach is mainly limited by the accuracy of the used FM and the used optimization scheme. As the mismatch-functional is non-convex and the dimensionality of the problem is too high for a global optimization, a local search faces the challenge of terminating in a local minimum. These minima are often sufficiently close to the global minimum, so that the particle parameters are approximated with great accuracy, as shown in the present work. With an increased number of particles in the measurement volume, the higher dimensionality introduces more local minima. This is a crucial insight, as the accuracy at which a PI is reconstructed does not depend of the PI overlap, but is rather subject to which local minimum the solver terminates in. This also shows the importance of good initialization. For a different initialization, therefore, a slightly different result is expected. This can be used to run an ensemble method (multiple optimization runs with random perturbations in initialization) on the same image to provide a measure of uncertainty for real experimental approaches.

The presented IP approach poses a promising method to refine detections from machine learning methods as conceptualized in the hybrid approaches [SDLK23]. While the IP approach results in larger computational cost compared to FFT approaches for sizing or a Hough transform for the position determination, the method may be used for off-line measurements to improve accuracy, where the time to results is not of the essence. The next development step for the IP approach would be to use the approach for two-phase flows, as different kinds of particles such as single emitters (e.g., tracers) and double emitters (bubbles/droplets) can easily be represented within the present approach (i.e., by omitting the diameter parameter for tracers). The capability of the IP approach to deal with low SNRs and tolerate aliasing when imaging particles closer to the focal plane to obtain better SNR, can improve the evaluation of data with poor quality. In consequence, this presents a promising step towards IPI in backscatter, which is suffering from poor data quality. Future steps could also include more complex particles such as multi-emitters like those arising from rough particles [23] or non-circular bubbles and droplets [188, 189]."[SGK25]

5 Simultaneous Measurement of two Phases in Defocusing Particle Tracking Velocimetry

This chapter addresses research question RQ4, on the challenge of enabling simultaneous two-phase measurements in DPTV, by exploring whether CNNs can be trained to detect PIs and classify them by phase. Building on prior work in neural network-based PI detection for single-phase DPTV [35, 56, 99, 14, 48, SDLK23, 223, 159], the focus in this chapter is on distinguishing between tracer particles in the continuous phase and dispersed-phase particles such as droplets or bubbles. This capability is essential for extending DPTV to DTPFs, where phase-specific tracking or IP approaches require reliable phase distinction. The content of this chapter is based on the manuscript *Pattern-Based Phase-Separation of Tracer and Dispersed Phase Particles in Two-Phase Defocusing Particle Tracking Velocimetry* by Sax and Kriegseis [SK25b].

5.1 Introduction

"Several methods have been developed to measure both phases simultaneously, often by combining different techniques." [SK25b] This chapter, however, focuses on measuring both phases using only the DPTV method. Before selecting a specific approach for phase distinction, it is useful to first consider how similar challenges have been addressed in PIV and PTV, where the extension to two-phase flows has been extensively studied; see, for example, [29, 97, 94, 31, 82]. A central concept in these extensions is the reliable differentiation between particle types, typically tracers and dispersed-phase particles such as bubbles or droplets. In essence, two main strategies are employed to achieve this distinction:

"The first approach employs wavelength-based techniques to distinguish between particle types. Dual-camera setups utilize wavelength filters to differentiate between particles with or without fluorescence or with fluorescence at different wavelengths [29, 100]. Alternative methods exploit the separate color channels of a single RGB camera for phase distinction [97]. Other techniques apply laser-induced fluorescence (LIF) to identify, for example, bubbles by detecting the absence of tracers in the carrier fluid. Wavelength-based techniques are advantageous due to their reliable phase distinction. However, multi-camera systems are complex and challenging to calibrate, and RGB cameras offer only lower sensor resolution.

The second category of approaches distinguishes the phases solely during data post-processing. These methods include exploiting size differences between particle types [94, 31, 82, 91], combining differences in particle size and the intensity of scattered light [91], and applying ensemble correlation of particle motion [43, 177] to separate the phases. Post-processing-based distinction generally permits the use of standard equipment, but depends on the classification accuracy of the underlying method. Ensemble-based techniques require a sufficient number of particles within the image, and size-based distinction is only feasible when the two particle types differ significantly in size.

Since DTPFs are inherently three-dimensional, their investigation should ideally be based on volumetric measurement techniques. While conventional two-phase PIV is typically limited to planar measurements, the same phase separation strategies can, in principle, be extended to three-dimensional PTV methods such as tomographic PTV [143, 126, 173, 174], enabling full 3D-3C (three dimensional - three components) Lagrangian velocity measurements. In applications with limited optical access, single-camera techniques such as defocusing PTV (DPTV) [215, 146] and astigmatism PTV (APTV) [33, 86] offer volumetric measurement capabilities. These methods determine the out-of-plane position of particles based on the degree of defocus in the particle image (PI). The present work focuses on the phase distinction in DPTV. However, the phase distinction in DPTV is less trivial: Although phase separation by wavelength is feasible in APTV and DPTV, it undermines the principal advantage of these methods being single-camera systems. Moreover, the use of RGB cameras can be problematic due to their lower sensor resolution, which reduces the accuracy of PI diameter estimation. Post-processing-based phase separation presents an alternative, however, distinguishing particles by size is particularly challenging in defocused images, as PI size is primarily governed by the degree of defocus rather than the actual particle size. Ensemble-based methods are likewise unsuitable for DPTV, as they require high seeding densities, which are often unfeasible due to the method's inherent limitation to lower seeding levels caused by PI overlap.

The present work proposes a post-processing approach for two-phase DPTV that employs neural networks to classify PIs based on their visual patterns. The method enables 3D position determination of both tracers and bubbles or droplets, using a single camera by exploiting differences in the appearance of defocused PIs, which result from distinct scattering behaviors. By relying on scattering characteristics rather than particle size, the approach permits tracer and dispersed phase particles to be of identical size and eliminates the need for specialized equipment such as RGB cameras, which may impose other limitations. As with any post-processing method, the effectiveness of the approach depends on the classification accuracy and the algorithm's ability to generalize across different particle types.

For the task of pattern-based phase distinction, the present approach employs convolutional neural networks (CNNs) [109, 111], which are increasingly adopted within the DPTV and APTV communities and have demonstrated enhanced performance in PI detection [35, 56, 99, 14, 48, SDLK23, 223, 159]. Modern object detection networks, such as Faster R-CNN [164] and YOLO [163, 161, 162], are capable not only of detecting but also classifying objects, rendering them a natural choice for distinguishing between different PI types. This work utilizes object detection networks to detect and classify different PI types based on their visual patterns. While this distinction is theoretically applicable to various particle types, the present study focuses exclusively on differentiating between tracer particles and droplets or bubbles. Consequently, the use cases addressed pertain specifically to gas-liquid DTPFs.

Nevertheless, the application of CNNs for pattern-based phase distinction in this context introduces several non-trivial challenges. While CNNs are effective for PI detection in DPTV, their performance is typically highest when the test data closely resembles the training data. However, variations in optical setups can significantly alter the appearance of PIs, thereby challenging the network's ability to generalize across different conditions and potentially reducing both detection and classification accuracy. To enhance generalization, the training dataset should be large and diverse, incorporating images from a variety of optical configurations. However, object detection networks typically require labeled data to be trained. The acquisition of such labeled data remains one of the principal barriers to the broader adoption of CNNs in DPTV. Alternatively, training data tailored to a specific experiment can be used to train a CNN on case-specific data, thereby reducing the need for generalization. However, this approach necessitates a method for acquiring and labeling experiment-specific training data. To address the problem of acquiring labeled

training data for CNNs in DPTV, the present work proposes a framework for generating large volumes of auto-labeled training data using generative adversarial networks (GANs) [67]. To evaluate the feasibility of CNN-based phase separation in dispersed two-phase DPTV, Faster R-CNN and YOLOv4 [17] are employed. These networks are assessed based on their ability to detect and classify PIs under both familiar and previously unseen imaging conditions."[SK25b]

In summary, the present chapter "introduces a method for the pattern-based distinction of tracers from bubbles or droplets using CNNs in two-phase DPTV, while concurrently addressing the principal challenge of acquiring suitable training data for these networks. This approach may also be extended by integrating interferometric particle imaging (IPI) [98, 64, 144], enabling simultaneous determination of the particle size of the dispersed phase from fringe patterns in the PIs."[SK25b]

5.2 Theoretical Concept of Pattern based Phase Distinction in Defocused Images

"The present two-phase approach for DPTV utilities the pattern of a PI to differentiate between tracers and bubbles or droplets. While defocused PIs of tracers and those of bubbles or droplets are fundamentally similar, the key distinction lies in the presence of unidirectional interference fringes in the latter, as typically observed in IPI. In contrast, PIs from opaque tracer particles, as used in DPTV, generally do not exhibit such fringe patterns. It should be noted, however, that certain transparent tracers may also produce fringe patterns, this issue is addressed later in the present work. These differing PI patterns arise from the distinct scattering characteristics of the particles. These characteristics, in turn, result from differences in particle surface properties and transparency, as illustrated in Fig. 5.1. Fig. 5.1 illustrates the three types of PI patterns along with their respective scattering characteristics.

For example, an opaque tracer may possess either a smooth or a rough surface, each influencing the scattering behavior differently. In the case of a smooth opaque tracer, light is reflected towards the camera at only a single point (glare point), resulting in a PI that exhibits no interference pattern, referred to as an *empty pattern*; see the left column of Fig. 5.1. Similarly, smooth, transparent tracers also exhibit only a single glare point under specific illumination angles. An example of such tracers is Di-Ethyl-Hexyl Sebacat (DEHS), which typically produces empty pattern PIs.

For tracers with a rough surface, light is scattered at multiple points across the surface, producing glare points without a geometrical order. These rough particles exhibit a speckle pattern [39], which appears as a random, non-directional pattern (*irregular pattern*); see the right column of Fig. 5.1. Polyamid tracers are an example of particles that produce such speckle patterns.

In contrast, bubbles or droplets, which are transparent particles with smooth interfaces, reflect and refract light at multiple points. This results in more than one glare point on the particle surface [201]. Notably, these glare points lie along a single line defined by the scattering plane, producing unidirectional interference fringes in the PI (*regular pattern*); see the middle column of Fig. 5.1. Thus, PIs of bubbles or droplets can be identified by their regular, unidirectional fringe patterns.

It should be noted that certain transparent tracers (e.g. DEHS tracers) may also possess smooth surfaces and can therefore exhibit unidirectional interference patterns, making them difficult to distinguish from bubbles or droplets. In such cases, alternative tracer materials should be selected. Occasionally, a smooth transparent tracer with a refractive index different from that of the bubbles or droplets may still produce an empty pattern PI. This occurs under specific scattering angles where only a single glare point is formed, while other particles of different refractive index may

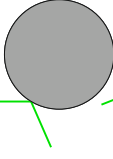
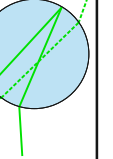
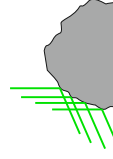
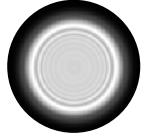
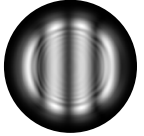



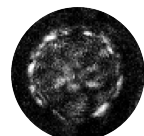
Particle Type	smooth opaque smooth transparent*	smooth transparent	rough opaque
Scattering Characteristic			
Simulated Particle Image	 ($n=1$ glare point)	 ($n=2$ glare points)	 ($n=12$ glare points)
Real Particle Image	 Di-Ethyl-Hexyl Sebacat (DEHS) tracer	 water droplet	 Polyamid tracer

Figure 5.1: "Three different types of PIs: PIs with an empty pattern (left), PIs with a regular unidirectional pattern (middle) and PIs with an irregular speckle pattern (right). The particle type, the scattering characteristic and simulated and real examples are shown. The simulated examples were created using the forward model from Sax *et al.* [SGK25] and the glare point based scattering characteristic. For the rough opaque particle a number of glare points significantly larger than two (here 12) were chosen to model multiple randomly distributed glare points. The real examples show a DEHS tracer from Exp. G2 (left), a water droplet from Exp. T4 (middle) and a Polyamid tracer from Exp. G1 (right)."[SK25b]. Figure from [SK25b].

exhibit two. However, this scenario is complex and requires careful consideration of the scattering characteristics of transparent particles.

Within the scope of this work, only gas-liquid DTPFs, such as bubbly flow, are considered. Consequently, neural networks were trained to distinguish between bubbles or droplets, characterized by regular, unidirectional fringe patterns, and tracers, which exhibit either empty or irregular speckle patterns. Accordingly, PIs with empty and speckle patterns were grouped into a single class (tracers), while PIs from bubbles and droplets were assigned to the other class (dispersed phase). The approach may also be applicable to other flow types, however, these lie beyond the scope of the present work and remain a subject for future investigation."[SK25b]

5.3 Auto-labeling Approach and Training of Networks for Particle Detection and Classification

"To employ CNNs for the detection and classification of different types of PIs, training data must first be acquired and labeled. In the context of DPTV, at least three distinct approaches have been reported in the literature for obtaining such training data: Images can be taken from calibration measurements to train the CNN [35]. However, labeling experimental data is time-consuming, since there is no known ground truth and typically results in suboptimal dataset sizes. An alternative is training on synthetic (model based) data [56, SDLK23]. Using physical models to generate synthetic images with known ground truth allows for automatic labeling and the creation of sufficiently large training datasets. However, synthetic images do not necessarily

represent real images with sufficient fidelity and often lack diversity in features, which may result in poor generalization capability of the network. Additionally, large dataset sizes are only partially beneficial if the variety in features is limited and the images become repetitive. A compromise between the two approaches is the use of image-generating deep learning methods, such as GANs, to generate synthetic (data based) images that more closely resemble real PIs [48].

An approach based on GAN-generated training data is proposed for object detection in two-phase flow experiments. An overview of the methodology is provided in Fig. 5.2. The objective is to train a CNN on data that closely replicates the conditions of the actual experiment. First, two data collection experiments are conducted using the same experimental setup as the target experiment (right side of Fig. 5.2). Each experiment captures only a single phase, either tracers or the dispersed phase. PIs are extracted from the resulting image sets using conventional algorithms. These extracted PIs are then used to train two GANs, each corresponding to a specific PI type (e.g., tracer or bubble). These GANs are then employed to generate synthetic images containing both PI types (upper part of Fig. 5.2). The synthetic mixed-phase images are used to train the object detection CNN. Once trained, the CNN is applied to the actual two-phase experimental data to detect and distinguish between different PI types (lower part of Fig. 5.2).




Figure 5.2: "Auto-labeling approach and PI detection. Two data acquisition experiments are conducted, each involving a single phase. PIs are extracted from the resulting image sets using conventional algorithms. These extracted PIs are then used to train GANs, to generate mixed-phase image sets. The mixed-phase sets are subsequently used to train an object detection CNN. Finally, the trained CNN is applied to the actual two-phase experimental data." [SK25b]. Figure from [SK25b]

This approach enables the generation of training data tailored to a specific experimental setup, as illustrated in Fig. 5.2. Alternatively, multiple data collection experiments can be performed to produce more diverse training data, thereby facilitating the development of object detectors capable of generalizing across various experimental configurations.

The approach can be divided into two steps: the first part involves the acquisition of raw images and the training of the GANs for the auto-labeling approach (upper part of Fig. 5.2), which is

presented in Section 5.3.2. The second part concerns the training of the actual object detection network used for phase distinction (lower part of Fig. 5.2), which is described in Section 5.3.3.

However, before the training and data acquisition procedures are discussed, the method for generating training data using GANs is elaborated in more detail in Section 5.3.1 (auto-label approach - right upper part of Fig. 5.2)."[SK25b]

5.3.1 Conceptual Approach for the Generation of Auto-labeled Training Data

"This section provides more detail on the auto-labeling approach using GANs, shown in Fig. 5.3. The auto-labeling approach employs two GANs trained on PIs of either tracers or bubbles, obtained from dedicated data acquisition experiments. The method operates by using a GAN for each phase to generate a PI, which is then inserted into a full image. Two GANs, one per PI type, are used so that the class of each generated PI is always known. This also offers the advantage that each GAN only needs to specialize in the features of a single PI type, avoiding the risk of feature mixing between classes.




Figure 5.3: "Generation of automatically labeled training data for object detection networks. Two GANs - one for tracer particles and one for the dispersed phase - independently generate small image snippets, each containing a single PI. For each snippet, a BB is placed over the PI. The image snippets and corresponding BBs are then randomly resized and inserted into an empty image."[SK25b]. Figure from [SK25b].

To train the GANs, single-phase data acquisition experiments are conducted to collect a large number of images containing only one PI type. These experiments must be single-phase so that the class of the PI in each image is known, and classification can be avoided. A conventional algorithm, such as the circular Hough transform [9, 7, 8], is then used to extract individual PIs from the images, see the Hough transform example in Fig. 5.2. The Hough transform also determines the size of the PIs, allowing them to be rescaled to a standard size, usable by the GAN. It should be noted that the data acquisition experiments should be conducted at low seeding densities to ensure reliable PI detection by the Hough transform, minimising false positives and avoiding PI overlap. PI overlap is further mitigated by calculating the intersection over union (IoU) for each detection and removing overlapping PIs from the dataset. In this way, a large number of individual PI examples can be collected for each PI type. While approach requires the recording of typically 100 to 1000 images per class, it eliminates the need for manual labeling.

Using these two PI sets, one for tracers and one for bubbles or droplets, the GANs can be trained to generate images containing PIs that resemble those in the recorded datasets. Once trained, the GANs generate image snippets, each containing a single PI of the corresponding type. This is achieved by also providing the GANs only with image snippets of individual PIs in training. For each GAN-generated image, the Hough transform is applied again to place a bounding box (BB) around the PI. The class label is automatically assigned based on the GAN that generated the image (i.e. tracer-GAN or bubble-GAN); see the BB placement step in Fig. 5.3.

The labeled PIs are then randomly resized using linear interpolation (up and down sampling) and placed at random positions within the image; see the red arrows in Fig. 5.3. The position and size of the PI BB are updated during the placement and resizing process to reflect their location in the full image. In this way, the BB position is known in the final image. Before inserting the PI into the image, a circular mask is applied to the GAN-generated image snippet outside the BB to subtract the background. This step reduces artifacts caused by pasting the snippet into the image and prevents the network from learning these artifacts as cues for PI detection. However, there remains a risk that the network may learn the mask itself as an unintended feature.

Using this method with two GANs, a random number of PIs per phase is placed into each image, with every PI assigned a BB and class label. This process can be repeated across multiple images to generate large training datasets that resemble those obtained from the data acquisition experiments. The advantage of using GANs over directly employing unlabeled PIs from the experiment lies in their ability to produce a greater number of images than the original dataset used for training. In this way, the GANs act as a multiplicative factor, effectively increasing the size of the dataset."[SK25b]

5.3.2 Training the Generative Adversarial Networks for Image Generation

"To realize this approach, raw images must be collected from data acquisition experiments to train the GAN networks. This section outlines the data acquisition procedures employed in this work, as well as the subsequent training of the GANs. While the proposed method can be used to tailor training data to a specific experimental setup, the resulting insights on the performance would be highly case-specific. Instead, a second strategy is adopted: generating a larger and more diverse training dataset by incorporating PIs from multiple experiments. This enables the CNNs to be trained for applicability across different experimental conditions and allows their generalization capability to be evaluated. As a result, the findings presented in this work are more transferable and relevant to a broader range of use cases."[SK25b]

Training Data Acquisition

"In this work, training images were collected from four different experiments: two containing only tracers (Exp. G1 and G2) and two containing only bubbles (G3 and G4). To create a large dataset with a diverse representation of PIs, only the visual appearance of the PIs is relevant. The underlying experimental conditions and hardware used in the individual acquisition experiments are not critical to this approach. However, a brief overview of the acquisition experiments is provided below:

1. Experiment G1 (*Tracer Particles*): Polyamid tracers (particle size $d_P = 20\text{ }\mu\text{m}$, PI size $d_{PI} = [19, 273]\text{ px}$) in a rotating flow within a glass cylinder. A total of 22,804 PIs, like the ones shown in Fig. 5.4a, were sampled. The experimental setup of the experiment is described in Lange *et al.* [LSBK24].

2. Experiment G2 (*Tracer Particles*): DEHS tracers (PI size $d_{PI} = [15, 22]$ px) in an airflow behind a plasma actuator. A total of 11,649 PIs, like the ones shown in Fig. 5.4b, were sampled. The images were sourced from the experiment described by Pasch *et al.* [149].
3. Experiment G3 (*Bubbles*): Bubble column in a rectangular glass tank (bubble size $d_P = [20, 250]$ μm , PI size $d_{PI} = [110, 122]$ px) . A total of 11,875 PIs, like the ones shown in Fig. 5.4c, were sampled. The images were sourced from the experiment described in Sax *et al.* [SDK25].
4. Experiment G4 (*Bubbles*): The same setup as Exp. G1 was used, but bubbles (particle size $d_P = [20, 300]$ μm , PI size $d_{PI} = [19, 273]$ px) were recorded instead of tracers. A total of 6,721 PIs, like the ones shown in Fig. 5.4d, were sampled.

Further details on the experiments can be found in the referenced literature [149, SDK25, LSBK24].




Figure 5.4: "PIs sampled from the Exp. G1[LSBK24] (a), Exp. G2[149] (b), Exp. G3[SDK25] (c) and Exp. G4[LSBK24] (d). These PIs were used to train the GANs for tracer particles (a)&(b) and for particles of the dispersed phase (c)&(d)."[SK25b]. Figure from [SK25b].

The two training datasets (Exp. G1 and G2) contained a total of 34,453 examples of tracer PIs used to train the tracer GAN, while 18,596 examples of bubble PIs from Exp. G3 and G4 were used to train the dispersed phase GAN."[SK25b]

Training of the Generative Adversarial Networks

"The GANs for tracers and dispersed phase particles share the same architecture and differ only in the training datasets used. MATLAB's framework and GAN architecture were employed for this purpose [131]. Details of the generator and discriminator architectures are provided in Tab. B.1

in Appendix B. Both GANs were trained for 50 epochs. The training specifications are listed in Tab. B.2, and the generator and discriminator scores during training are shown in Fig. B.1 in Appendix B.

The GAN for the dispersed phase, which was trained on the smaller of the two datasets, exhibited a divergence between the generator and discriminator scores beyond the 29th epoch, as shown in Fig. B.1. This behaviour, indicative of degradation, led to the selection of the checkpoint from the 29th epoch for further use. Notably, the discriminator score approached 0.5 during training, indicating that the discriminator became increasingly uncertain in distinguishing real from generated data—an expected sign that the generator was producing more realistic outputs. The resulting PIs generated by both GANs are shown as examples in Fig. 5.5a for tracers and in Fig. 5.5b for dispersed phase PIs."[SK25b]




Figure 5.5: "GAN-generated images of PIs of tracer particles (a) and particles of the dispersed phase (b). All images have dimensions of 64×64 px."[SK25b]. Figure from [SK25b].

5.3.3 Training of Object Detection Networks

"Using the outlined approach to generate auto-labeled training data, 20,000 images containing half a million PIs were created. These images were augmented through random vertical and horizontal flipping, resulting in a dataset of approximately 2 million PI examples, equally distributed between the two classes."[SK25b]

Selection of Object Detection Networks

"Two different object detection architectures were evaluated to assess the feasibility of two-phase DPTV using CNNs. However, these architectures serve as representative examples, and other models may be used as well. Faster R-CNN [164] and YOLOv4 [17] were selected to include both a two-stage and a single-stage object detector, respectively. These two networks were chosen, as both have previously been employed for single phase DPTV [35, 99, 48, SDLK23]. The Faster R-CNN network, using a ResNet50 backbone [72] (42 million parameters), was trained with the TensorFlow framework [1]. This architecture has demonstrated strong performance in both DPTV and APTV applications [35, 48, SDLK23]. While YOLOv3 [162] has shown inferior results compared to Faster R-CNN in DPTV tasks [48], YOLOv4 [17] incorporates architectural improvements that makes it a promising candidate. YOLOv4 introduces spatial pyramid pooling (SPP) [71], which enhances the network's ability to detect objects at multiple scales, particularly smaller ones. Additionally, YOLOv4 replaces the multi-scale detection approach of YOLOv3 with a more

advanced path aggregation network (PANet) [122], improving feature fusion across scales. These enhancements result in finer-resolved feature maps with higher semantic information, which were hypothesized to improve PI detection in DPTV [48]. YOLOv4 also employs cross stage partial (CSP) connections [206], which improves generalization capability. To investigate whether large backbones are necessary for DPTV, two versions of YOLOv4 were trained. The first used a large CSP-DarkNet-53 backbone, a combination of DarkNet-53 [162] and CSP-Net [206], with 63.9 million parameters. The second used a significantly smaller CSP-DarkNet-19-tiny backbone with 5.9 million parameters for comparison."[SK25b]

Configuration and Training Process

"The training was initialized from weights pre-trained on MS-COCO 2018 [118] for both networks, to leverage generic features. In this transfer learning approach [37], the networks only need to learn domain-specific features for detecting PIs in each phase of DPTV, which typically results in shorter training times. The training parameters are detailed in Tab.B.3 in the Appendix B. The training dataset was split (prior to augmentation) into 90% for training and 10% withheld for the evaluation of the validation loss. Hyperparameters were selected based on the recommended settings used during pre-training on MS-COCO.

Faster R-CNN was configured to train for 20,000 iterations without early stopping, based on insights from previous work using this architecture for DPTV [SDLK23]. Both YOLOv4 variants were set up with a maximum of 20,000 training iterations. However, since tiny YOLOv4 is expected to converge faster than the full YOLOv4 network, early stopping with a patience of 50 validation steps was incorporated to prevent overfitting. During training, the validation loss was evaluated every 50 iterations, and training was automatically terminated if no improvement in validation loss was observed for 50 consecutive evaluations. As a result, training of tiny YOLOv4 was terminated early at 10,850 iterations, with the final model selected based on the checkpoint with the lowest validation loss. The divergence between training and validation loss suggested overfitting. As a cross check, a second version of tiny YOLOv4 was trained from the 10,850-iteration checkpoint to the full 20,000 iterations. This version performed worse on all test datasets, confirming the suspicion of overfitting, and is therefore not further considered. The larger YOLOv4 model was also trained with the same early stopping strategy but completed the full 20,000 iterations without triggering early termination.

The training losses of all three networks are shown in Fig. B.2 in the Appendix B. For both YOLOv4 and Faster R-CNN, the total loss did not decrease significantly after 15,000 iterations, suggesting that training had converged."[SK25b]

5.4 Detection and Separation of Different Particle Image Types

"To evaluate the performance of the object detection networks, images from various experiments were sampled and manually labeled."[SK25b]

5.4.1 Test Dataset

"Similar to the training dataset, two datasets containing only tracers and two datasets containing only bubbles or droplets were used. For tracers, the test dataset T1 (see Fig. 5.6a) contains tracer PIs similar to those in G1, while the image set T2 (see Fig. 5.6b) differs significantly from the

training data. For dispersed phase PIs, the test dataset T3 (see Fig. 5.6c) contains bubble PIs similar to those in G3, while the image set T4 (see Fig. 5.6d) contains droplet PIs that differ significantly from the training data. This was done to test on both familiar and unfamiliar data.

The advantage of using images containing only a single phase is that the ground truth class of each PI is known with absolute certainty. This would not be the case in real images with two types of PIs, where human error in manual labeling could distort the evaluation. Since the object detector was not informed that only one type of PI was present in each image, the classification performance remains valid. The detector processes each instance independently, without relying on the overall image composition, so the absence of a second class does not affect the validity of the results.

The object detectors were also tested on synthetic (Exp. T5, see Fig. 5.6e) and real experimental images (Exp. T6, see Fig. 5.6f) containing two types of PIs. For the synthetic images, the ground truth class labels are known, allowing classification performance to be evaluated. In contrast, in the real mixed-phase dataset (T6), BBs could be drawn; however, the manual class labeling of PIs proved unreliable. To avoid drawing conclusions based on potentially flawed ground truth, only the detection but not the classification performance was assessed for this dataset.




Figure 5.6: "Exemplary images from the six test datasets: T1[LSBK24] (a) and T2[115] (b) show tracer particles in a single phase flow. T3[SDK25] (c) and T4[151] (d) show only bubbles and droplets but no tracers. T5 (e) shows a synthetic image of a two phase flow of bubbles and tracers and T6[LSBK24] (f) show a real image of bubbles and tracers."[SK25b]. Figure from [SK25b].

Images for the test dataset were drawn from the following experiments and synthetic data generation:

1. Experiment T1 (*Tracer Particles*): This experiment used the same setup as Exp. G1, but the images were withheld from GAN training. A total of 773 PIs was manually labeled. An

example image is shown in Fig. 5.6a. The experimental setup of the experiment is described in Lange *et al.* [LSBK24].

2. Experiment T2 (*Tracer Particles*): Tracers (particle size $d_P = 9.84 \mu\text{m}$, PI size $d_{PI} = [15, 31] \text{ px}$) in an open wet clutch. More details can be found in Leister *et al.* [115]. A total of 394 PIs was manually labeled. An example image is shown in Fig. 5.6b.
3. Experiment T3 (*Bubbles*): This experiment used the same setup as Exp. G3, but the images were withheld from GAN training. More details can be found in Sax *et al.* [SDK25]. A total of 202 PIs was manually labeled. An example image is shown in Fig. 5.6c.
4. Experiment T4 (*Droplets*): Water droplets (particle size $d_P = [20, 60] \mu\text{m}$, PI size $d_{PI} = [16, 30] \text{ px}$) in a spray downstream of a nozzle. The images were sourced from the experiment described in Pöppe [151]. A total of 554 PIs was manually labeled. An example image is shown in Fig. 5.6d.
5. Experiment T5 (*Two-phase - synthetic*): Synthetic image of bubbles (simulated particle size $d_P = [20, 250] \mu\text{m}$, PI size $d_{PI} = [43, 112] \text{ px}$) and tracers (PI size $d_{PI} = [43, 112] \text{ px}$). The forward model described in Sax *et al.* [SGK25] was used with the experimental settings from Exp. G3 (see [SDK25]) to generate the images. A total of 3,000 PIs was labeled. An example image is shown in Fig. 5.6e.
6. Experiment T6 (*Two-phase - real*): Tracers and bubbles in the same experimental setups as Exp. G1 and G4 (same tracer and bubble sizes, same PI sizes). A total of 304 PIs was manually labeled. An example image is shown in Fig. 5.6f. No class labels were taken from this dataset.

The classification performance was rigorously evaluated on synthetic and real single-phase datasets, where class labels are known with absolute certainty. The synthetic dataset consists of idealized PIs that differ only in the presence or absence of fringe patterns, making it particularly well-suited to assess whether the networks have learned this key distinguishing feature. Importantly, the object detector was also tested on the real mixed-phase dataset to evaluate detection performance, ensuring that its ability to generalize to realistic conditions was assessed." [SK25b]

5.4.2 Evaluation Metrics

"The evaluation of the neural networks for two-phase DPTV is divided into two parts, assessing detection and classification performance separately. This separation enables a clearer understanding of the network's ability to distinguish between phases in two-phase DPTV, independent from the influence of PI detection.

The detection performance is evaluated in a class-agnostic manner. A detection is counted as a true positive (TP_D) if the predicted BB has an IoU greater than 0.5 with a ground truth (GT) BB, regardless of the predicted class (class-agnostic). This IoU threshold is chosen for comparability with previous work [48, SDLK23]. A class-agnostic confusion matrix is computed: detections without a corresponding GT BB are counted as false positives (FP_D), and missed GT BBs are counted as false negatives (FN_D). Detection performance is typically evaluated using precision P

$$Pr = \frac{TP_D}{TP_D + FP_D}, \quad (5.1)$$

which represents the probability that a prediction is correct, and recall \mathcal{R}

$$\mathcal{R} = \frac{\text{TP}_D}{\text{TP}_D + \text{FN}_D}, \quad (5.2)$$

which is the fraction of detected PIs. However, both precision and recall depend on the chosen confidence score threshold, which reflects the network's confidence in a given prediction. To eliminate this dependency, the average precision (AP) is computed from precision and recall values across confidence thresholds ranging from 0.001 to 0.999, using the trapezoidal rule:

$$\text{AP} = \int_0^1 \text{Pr}(\mathcal{R}) d\mathcal{R} \quad (5.3)$$

The class-agnostic AP provides a comprehensive, threshold-independent metric for evaluating detection performance. However, AP penalizes the network if it fails to achieve high recall, even when precision is high. Since a detector can still be reliable within lower recall ranges and some missed detections may be acceptable, the truncated AP (TAP) is also computed and normalized by the max recall:

$$\text{norm TAP} = \frac{1}{\max(\mathcal{R})} \int_0^{\max(\mathcal{R})} \text{Pr}(\mathcal{R}) d\mathcal{R}, \quad (5.4)$$

which measures class-agnostic detection performance up to the maximum achieved recall. As TAP depends on the maximum recall reached, both TAP and the maximum recall are considered together.

To investigate the classification performance independently from detection, a class confusion matrix is computed based on all TP_D detections. No confusion matrix is computed for FP detections, as there is no GT class; for FN detections, since no class was predicted; or for true negatives (TN), as neither GT class nor predictions exist.

The detector's ability to classify a detection correctly is then evaluated using the classification accuracy

$$\text{CA} = \frac{\text{TP}_C + \text{TN}_C}{\text{total number of TP}_D}. \quad (5.5)$$

This represents the fraction of correctly classified detections among all TP detections and can be interpreted as the probability that a TP detection is also classified correctly by the detector. Since the classification accuracy also depends on the confidence score threshold, it is evaluated at the threshold corresponding to the maximum F1-score

$$\text{F1-score} = 2 \frac{\text{Pr} \times \mathcal{R}}{\text{Pr} + \mathcal{R}} \quad (5.6)$$

which identifies the point at which precision and recall are optimally balanced. The classification accuracy only provides insight into the classification of TP detections, but a detector may also produce FP detections. Since there is no GT class to evaluate the classification of FP predictions, the class bias is computed for FP_D predictions instead of the accuracy via

$$\text{Class Bias} = \frac{\text{FP}_{D,\text{Class1}} - \text{FP}_{D,\text{Class2}}}{\text{FP}_{D,\text{total}}}. \quad (5.7)$$

The class bias takes a value of +1 if all FP detections are predicted as tracers, and -1 if all are predicted as dispersed phase, and is zero if the FPs are equally balanced between classes. This provides insight into whether the detector is biased toward predicting one class over the other. Class bias is also evaluated at the threshold corresponding to the maximum F1-score.

Since the localization accuracy of neural network detectors for DPTV and APTV has been extensively studied in previous works (see, e.g., [35, 99, 14, 48, SDLK23, 159]), and the primary focus of this "[SK25b] chapter "is the distinction between tracers and bubbles or droplets, no further evaluation of localization performance is conducted in this work." [SK25b]

5.4.3 Detection Results

"The results of the networks detection performance is given in Tab. 5.1. The networks achieve high average precision across all test sets. Particularly strong performance is observed on images that closely resemble the training data (Exp., T1 and T3). As expected, average precision is lower on the test sets (T2 and T4) that diverge more substantially from the training images. This highlights the effect of domain shift between familiar and unfamiliar types of images on the CNN detection performance, and emphasizes the need for the networks to generalize more effectively. When generalizing to unfamiliar images, the primary limiting factor is the truncation of the maximum achievable recall, while precision remains high. This becomes evident from the norm TAP values, which are close to 1 in most test cases, indicating reliable detection performance within the recall limitations of the network, even on images requiring generalization. This high precision in images visually different from the training data, is an important insight, as high precision helps prevent false positives and, consequently, ghost particles in the tracking step, which could lead to erroneous velocity vectors. Preventing such ghost particles (i.e., achieving high precision) is typically more critical than achieving high recall.

Table 5.1: "Class agnostic PI detection performance for all six test sets. Shown are the average precision, the max recall and the normalized truncated average precision. The best score for each metric and each test set is highlighted in bold print." [SK25b]. Table from [SK25b].

Test Set	Metric	Faster R-CNN	tiny YOLOv4	YOLOv4
Exp. T1 <i>Tracers only</i>	Average precision	0.885	0.864	0.856
	max Recall	0.885	0.864	0.892
	norm TAP	1	1	0.960
Exp. T2 <i>Tracers only</i>	Average precision	0.667	0.713	0.729
	max Recall	0.667	0.714	0.729
	norm TAP	1	0.984	1
Exp. T3 <i>Bubbles only</i>	Average precision	0.952	0.989	0.957
	max Recall	0.952	0.989	0.957
	norm TAP	1	1	1
Exp. T4 <i>Droplets only</i>	Average precision	0.765	0.819	0.850
	max Recall	0.862	0.819	0.850
	norm TAP	0.888	1	1
Exp. T5 <i>Tracers & Bubbles Syn</i>	Average precision	0.997	0.993	0.993
	max Recall	0.997	0.993	0.993
	norm TAP	1	1	1
Exp. T6 <i>Tracers & Bubbles Real</i>	Average precision	0.891	0.736	0.826
	max Recall	0.891	0.900	0.916
	norm TAP	1	0.818	0.903

Following the detection, the next important step is the phase separation. The results for the PI classification are given in Tab. 5.2. Generally, high classification accuracy is observed, with values typically ranging between 95-100%. The exception is the performance of Faster R-CNN on Exp. T4 and YOLOv4 on Exps. T4 and T5, both of which are test sets that require greater generalization. Tiny YOLOv4, however, showed consistent classification performance across all test images, regardless of the similarity of their visual appearance to the training data.

Exp. T5 provides a particularly interesting case, as the images are nearly noise-free and the PIs are represented with idealized features, free from aberrations or other effects that the network might otherwise use as cues. At the same time, the visually distinct appearance of synthetic images to the training data, requires the network to generalize. The strong classification performance (class accuracy of 99.2% for Faster R-CNN and 95.5% for tiny YOLOv4) indicates that the networks have learned physically meaningful features, specifically, the ability to distinguish PIs from different phases based on the presence of fringes in the PI.

Table 5.2: "Classification results on the six different test datasets. The class accuracy for true positive detections and the class bias of false positives is given. The bias becomes +1 if all FP are predicted to be tracers and -1 if predicted as dispersed phase. The best score for each metric and each test set are highlighted in bold print." [SK25b]. Table from [SK25b].

Test Set	Metric	Faster R-CNN	tiny YOLOv4	YOLOv4
Exp. T1 <i>Tracers only</i>	Class Accuracy (TP)	0.971	0.961	0.970
	FP Class Bias	+0.653	+0.868	+0.639
Exp. T2 <i>Tracers only</i>	Class Accuracy (TP)	0.983	0.998	1
	FP Class Bias	+0.365	+0.994	+0.915
Exp. T3 <i>Bubbles only</i>	Class Accuracy (TP)	1	0.994	1
	FP Class Bias	-0.640	-0.468	-0.658
Exp. T4 <i>Droplets only</i>	Class Accuracy (TP)	0.754	0.966	0.814
	FP Class Bias	-0.136	-0.509	-0.163
Exp. T5 <i>Tracers & Bubbles Syn</i>	Class Accuracy (TP)	0.992	0.955	0.824
	FP Class Bias	-0.622	-0.307	-0.264
Exp. T6 <i>Tracers & Bubbles Real</i>	Class Accuracy (TP)	-	-	-
	FP Class Bias	-0.357	-0.229	-0.410

The previous results suggest a low number of false positive (FP) detections, indicated by the norm TAP values close to 1, at reasonable recall levels and reliable phase separation for true positive (TP) detections. However, the networks still exhibit class bias when making FP detections. It is important to investigate FP class bias, as false positives introduce non-existent particles into the tracking process, which can influence the accuracy of the determined phase-specific velocity vectors. In the single-phase images (Exp. T1 to T4), the FP class bias tends toward the class present in the image. This is primarily a result of duplicate detections within PI clusters in images where the class distribution is imbalanced. In contrast, Exps. T5 and T6 contain both classes in balanced proportions. In these mixed-case scenarios, the network tends to show an FP bias toward the dispersed phase, which increases with the confidence threshold (see Fig. B.3 in Appendix B). Since the training data was class-balanced, a likely explanation is that PIs of bubbles and droplets are more visually distinct than PIs from tracers, leading the network to be more confident in predicting this class.

The generalization capability to visually different data can be assessed based on the results of Exps. T2, T4, T5 and T6. It shows that Faster R-CNN generalized better to the synthetic data, indicating that it learned meaningful physical features, whereas YOLOv4 generalized more effectively to different real datasets, both in detection and classification."[SK25b]

5.5 Concluding Remarks

This chapter "demonstrates that pattern-based phase distinction using CNNs is a reliable and effective method for dispersed two-phase DPTV. By leveraging the distinct visual appearance of defocused PIs, the approach enables accurate classification of tracer particles versus bubbles or droplets, achieving 95-100% accuracy across diverse datasets. When combined with double- or multi-frame setups, this method allows for simultaneous and distinct 3D-3C tracking of both phases using only a single camera. Furthermore, the approach may be extended by incorporating IPI to enable sizing of bubbles or droplets present in the image.

A key finding is that the CNNs appear to have learned physically meaningful features to distinguish tracers from bubbles or droplets, rather than relying on superficial correlations. A major advantage of the proposed method is that, since the phase distinction is performed entirely in post-processing, it does not require any additional equipment, unlike wavelength-based separation techniques. Furthermore, being a single-camera approach, it enables 3D-3C two-phase tracking with only one optical access, making it particularly attractive for applications with limited optical accessibility.

Beyond its core effectiveness, the method offers several practical advantages. It is particularly well-suited to scenarios where specialized equipment or additional cameras are unavailable or impractical, thereby excluding the use of wavelength-based phase distinction. Unlike size-based separation techniques, the proposed method remains effective even when the size distributions of dispersed and tracer particles overlap. Furthermore, the pattern-based approach does not rely on large velocity differences or high seeding densities, which are typically required for ensemble correlation-based phase separation. However, as a post-processing-based approach, the method depends on the accuracy of the employed CNNs and may be less reliable in phase distinction compared to wavelength-based alternatives. Additionally, the approach entails increased computational demands relative to other post-processing-based phase separation methods.

A central challenge in applying CNN-based phase distinction in two-phase DPTV lies in the availability of sufficiently large and representative training datasets. To address this, the present work introduces a GAN-based auto-labeling framework that enables the generation of realistic, experiment-specific training data from raw, unlabeled images. This approach significantly reduces the manual effort typically required for data annotation and enhances the applicability of CNNs in experimental fluid mechanics. The method involves conducting single-phase calibration experiments using the same optical setup as the target measurement. The resulting raw images are then used to train a generative model capable of producing large volumes of labeled data that closely resemble the target application. This tailored dataset reduces the domain gap between training and deployment, thereby improving the detection and classification performance of the CNNs. Moreover, the framework can be extended to incorporate data from multiple experimental setups, facilitating the creation of more diverse training datasets. Such diversity enhances the generalization capability of the trained networks, making them more robust across varying experimental setups and reducing the need for retraining. The results presented in this work demonstrate that both Faster R-CNN and YOLOv4 exhibit improved performance on familiar data. While the auto-labeling approach introduces additional computational cost, requiring the training of both

a generative model and a detection network, it eliminates one of the primary bottlenecks in the adoption of CNNs for DPTV: the need for extensive manual labeling. Furthermore, the method is also applicable to single-phase DPTV, where only one GAN is required, further broadening its utility. In summary, the proposed auto-labeling framework constitutes a critical step toward scalable and adaptable CNN-based DPTV.

To improve the approach of pattern based phase distinction in two-phase DPTV, the following considerations should be made: First, while the CNNs demonstrated very high TAP with, this performance was achieved only within limited maximum recall ranges. Generalization to visually unfamiliar images had little impact on precision but significantly constrained recall. Although high precision is generally more critical in DPTV to avoid ghost particles and the resulting erroneous velocity vectors, higher recall remains desirable. Improved recall would increase the density of velocity vectors and reduce the risk of missing substantial portions of the numbers of particles, which could otherwise lead to incomplete flow characterization. Moreover, achieving recall values close to 100% for the dispersed phase, in combination with particle sizing via IPI, would enable void fraction estimation. Therefore, future efforts should prioritize increasing recall. This could be achieved either by minimizing the required generalization, through training data that more closely represents the target experiment, or by constructing larger and more diverse datasets to improve the network's generalization capability. Notably, YOLOv4 exhibited better generalization to real data and may be a promising candidate for future development. Second, attention should be given to mitigating the FP class bias observed in the CNNs, which tended to favor bubbles or droplets. Although the high precision achieved by the networks resulted in very few FPs, this bias could become significant when operating at higher recall levels. The presence of this bias, despite a class-balanced training dataset, suggests that additional measures are necessary. One potential strategy is to introduce a slight class imbalance in favor of tracer particles during training, thereby increasing the network's exposure to the class it is less confident in distinguishing. Finally, incorporating synthetic examples of PIs alongside real ones in the training data may prove beneficial. Exposure to idealized features during training could help the networks learn more physically meaningful representations and reduce reliance on context-based cues. These directions collectively aim to improve the robustness, generalization, and quantitative reliability of CNN-based phase distinction in two-phase DPTV."[SK25b]

In summary, this chapter "confirms that pattern-based phase distinction using CNNs is a viable and practical solution for two-phase DPTV, with promising avenues for further improvement through targeted training data and network design."[SK25b]

6 On the Particle Image Overlap in Single Camera Defocusing Approaches without Astigmatism

This chapter concerns the research question RQ5 on how PI overlap is affected by experimental settings such as the number of PIs in an image or the degree of defocusing. The content of this chapter builds upon the manuscript *On the Particle Image Overlap in Single Camera Defocusing Approaches* by Sax and Kriegseis [SK25a].

6.1 Introduction

"A major concern in both DPTV [215, 146] and IPI [98, 64] is particle image (PI) overlap. In DPTV, this PI overlap makes identifying individual PIs more difficult, as their boundaries are partially covered by other PIs, compare the left example in Fig. 6.1. This impairs the detection capability of detection algorithms [57, 14, 48, SDLK23]. PI overlap complicates the measurement of the defocused particle image diameter, d_{PI} . As a result, fewer particles can be imaged simultaneously to avoid overlap, which in turn reduces the spatial resolution of the reconstructed flow field. As the PI boundary is partially covered in PI overlap, the covered part of the boundary must be assumed based on the curvature of the uncovered part. This leads to an estimation of d_{PI} rather than a measurement of d_{PI} for some PIs (see middle part of Fig. 6.1), which in turn increases the uncertainty of the reconstructed z -position of the particle. Other approaches simply ignore PIs with overlap to avoid this issue [57]. Not being able to estimate the PI's boundaries can also affect the determination of the PI's center coordinates (x_{PI}, y_{PI}) , which can additionally impair subsequent in-plane velocity estimations, from a time series of images. The same effects, i.e. limited number of detectable particles and limited z -position accuracy, influence the evaluation of particles in IPI. However, in IPI, an additional factor comes into effect. As part of the PI area is covered, only a portion of the interference pattern can be used to size the underlying particle (d_P), compare the right example in Fig. 6.1. Consequently, multiple characteristics are of interest: In DPTV, the number of overlaps a PI experiences is of interest, as this provides an estimate of how much of the PI boundary will be covered. Simultaneously, the number of PIs in an image that do not experience PI overlap is also of interest, as it offers an estimate of how many particles might be reconstructed with greater accuracy or be used at all. Additionally, the maximum amount of overlap a PI experiences is of interest, as it provides information on the size of the boundary segments that are covered, which influences the determination of d_{PI} . For IPI, the primary metric of interest is the fraction of the PI area that is covered, or its complement, the remaining free PI area. The remaining assessable PI area is the portion of the PI that can still be used to determine the particle diameter d_P without the influence of PI overlap. These metrics are of interest, as evaluation and detection algorithms can be tested for the amount of overlap they can tolerate [14, 48, SDLK23].

However, these algorithmic limitations only provide indirect qualitative guidance for an experiment. They do not offer precise recommendations on how much defocusing can be tolerated or how big




Figure 6.1: "Three issues caused by PI overlap: The detection of covered PIs is more difficult (left). The PI size measurement becomes more challenging and the PI size might be only estimated due to the PI boundaries being covered (middle). The free remaining area (FRA) that can be used in IPI becomes small due to PI overlap (right)." [SK25a]. Figure from [SK25a].

the aperture diameter can be chosen (i.e. increase of d_{PI}), or how many particles N_P can be used in an experiment, to remain within the limits that the post-processing algorithms can handle reliably. It is of interest to use larger aperture diameters, as this allows more light to reach the camera chip and consequently, improves the signal-to-noise ratio. More importantly, larger aperture diameters increase the proportionality factor between the particle distance to the focal plane and the PI diameter, effectively enhancing the defocus sensitivity and thereby improving the depth position accuracy. Furthermore, greater defocusing can be advantageous for extending the measurement volume in the z -direction. However, both effects increase the PI diameter and, therefore, the likelihood of PI overlap.

Therefore, the objective of the present study is to provide guidance on the extent to which the PI size (i.e. the amount of defocusing and aperture diameter) and the number of PIs can be increased before PI overlap becomes problematic. Furthermore, the trade-off between PI size and source density for constant PI overlap is discussed. A statistical analysis of PI overlap as a function of PI size d_{PI} and the number of particles N_P is conducted. This analysis allows for the translation of algorithmic limitations into practical experimental settings. The theoretical investigation assumes circular PIs, neglecting the influence of astigmatism, and considers uniformly sized PIs. However, testing the model on four different experiments, including cases with non-uniformly sized PIs diameters and slight astigmatism, reveals that the model remains valid despite these simplifying assumptions." [SK25a]

6.2 Datasets for the Empirical Investigation

6.2.1 Generation of a Dataset for an Empirical Investigation

"To conduct a statistical investigation of PI overlap, a dataset of images containing PIs is required first. However, experimental images cannot be used for this purpose due to the previously discussed limitations of detecting PIs with high degrees of overlap. Instead, synthetic images are generated in which PIs are randomly distributed, allowing full control over and access to the ground truth positions and sizes of all PIs.

The positions of the PIs in the images were generated under the assumption of a uniform distribution of PI center coordinates, analogous to the commonly assumed uniform particle distribution in Particle Image Velocimetry (PIV). However, this assumption does not hold under all experimental conditions. In certain cases, such as sprays or bubble columns, the spatial distribution of PIs is influenced by the underlying flow structures and is, therefore, non-random. For the purposes of

the present study, the uniform distribution serves as a practical and simplifying approximation. Given the wide variety of possible flow topologies, assuming a uniform distribution, simplifies the analysis while still capturing the essential statistical behavior. As will be demonstrated later, the uniform distribution provides a sufficiently accurate approximation for the intended analysis.

To ensure realistic PI sizes relative to the field of view, the optical parameters from the bubble sizing experiment in Sax *et al.* [SDK25] were used to determine the PI diameters. This experiment involves a column of air bubbles ranging from 20-200 μm in diameter, generated by electrolysis and imaged out-of-focus within a 7×7 mm field of view in a rectangular water tank. Further details can be found in Sax *et al.* [SDK25]. However, all dimensional quantities were converted into a dimensionless form, making the results applicable to any optical setup. Both the size and number of PIs were varied in the synthetic images. To reduce the dimensionality of the parameter space and simplify the evaluation, uniformly sized PIs were assumed. This represents scenarios in which particles are located at a single z -position, such as in light sheet-based measurements. In cases involving non-uniformly sized PI diameters, the mean PI diameter can be used as a representative value for computing overlap metrics. Although the broader size distribution is expected to increase the variance of these overlap metrics, it is not expected to significantly affect their mean values.




Figure 6.2: "Overview over the dataset for the empirical investigation of the PI overlap visualized. Six different number of particles per image $N_P = [10, 20, 30, 40, 50, 60]$ and ten different defocus lengths (i.e. PI sizes) $z_P = [1, 1.5, 2, 3, 4, 5, 6, 7, 8, 9]$ mm were investigated, resulting in 6×10 sets of 400 images each. The images were generated using the physical model from Sax *et al.* [SGK25] with parameters from the experiment described in [SDK25]."[SK25a]. Modified figure from [SK25a]."

Six different numbers of PIs per image are considered: $N_P = \{10, 20, 30, 40, 50, 60\}$. The PI diameter is influenced by both the aperture diameter and the defocus length. While the aperture diameter affects the proportionality factor between defocus and PI diameter, defocusing itself shifts all PIs toward larger diameters. However, since both the aperture diameter and the z -position

(when sufficiently far from the focal plane) have a linear effect on PI size, only the z -position was varied in this study to control the PI diameter. In total, ten different defocus lengths are investigated: $z_P = \{1, 1.5, 2, 3, 4, 5, 6, 7, 8, 9\}$ mm. The total parameter space thus spans 10×6 subsets. The PI size corresponding to each defocus distance was calculated using the formula from Shen *et al.* [179], which relies on the linear defocus relationship valid under sufficient defocusing conditions [57]. Consequently, image generation was reduced to placing a fixed number of circular PIs, each with a computed defocus diameter, at random positions within the image. The structure of the test dataset is illustrated in Fig. 6.2. Each of the 10×6 subsets contains 400 images to ensure statistical convergence of the evaluated parameters. Convergence validation is shown in Fig. C.1 in the Appendix C."[SK25a]

6.2.2 Experimental Datasets for Validation

"In the synthetic dataset, two key assumptions were made: (1) PIs are uniformly distributed across the image domain, and (2) all PIs within a given image possess an identical diameter. Although these assumptions may not strictly hold under real experimental conditions, they provide sufficiently accurate approximations, as will be demonstrated through comparison of the theoretical model with the validation experiments. To evaluate whether the empirical model from the synthetic dataset generalizes to real-world scenarios, four validation image sets were used. These sets originate from four distinct experiments, each involving different flow topologies and optical setups. Representative images from each dataset are shown in Fig. 6.3.




Figure 6.3: "Exemplary images from the four different experiments used to validate the PI overlap model. , Exp. 1[115] shows tracers in a gap flow (a), Exp. 2[LSBK24] shows tracers and bubbles in a vortex like rotating flow (b) and Exp. 3[SDK25] shows a bubble column in a resting fluid (c), Exp. 4[151] shows droplets in a spray behind a nozzle (d)." [SK25a]. Figure from [SK25a].

The first experiment shows tracer particles in an open wet clutch, which is a gap flow between a rotating disk and a stationary wall. A total of 30 images, like the one in Fig. 6.3a, containing 772 PIs were manually labeled (tracer size $d_P = 9.84 \mu\text{m}$, PI diameter $d_{PI} = [15, 31] \text{ px}$). More details on the experiment can be found in Leister *et al.* [115]. The second experiment shows the PIs of both tracers (tracer size $d_P = 20 \mu\text{m}$, PI diameter $d_{PI} = [19, 273] \text{ px}$) and bubbles (bubble size $d_P = [20, 250] \mu\text{m}$, PI diameter $d_{PI} = [19, 273] \text{ px}$) in a glass cylinder (outer diameter 105 mm, wall thickness 5.8 mm). The presence of the glass cylinder introduced mild astigmatism in the captured images, with the maximum observed aspect ratio (width/height) reaching 1.66. A magnetic stirrer was used to create a vortex-like, rotating flow. A total of six images, like the one in Fig. 6.3b, containing 303 PIs were labeled manually. Further information on the experiment can be found in Lange *et al.* [LSBK24]. The third experiment shows a column of air bubbles (bubble size $d_P = [20, 250] \mu\text{m}$, PI diameter $d_{PI} = [110, 122] \text{ px}$) in an otherwise stationary fluid (water). A total of 40 images, like the one in Fig. 6.3c, containing 201 PIs were labeled manually. More details on the experiment can be found in Sax *et al.* [SDK25]. The fourth experiment involves spraying droplets (droplet size $d_P = [20, 60] \mu\text{m}$, PI diameter $d_P = [16, 30I] \text{ px}$) from a nozzle outlet. A total of five images, similar to Fig. 6.3d, with a total of 451 PIs, were labeled manually. Details on the experiment can be found in Pöppe [151]. These four experiments represent markedly different flow scenarios: Experiments 1 and 2 involve shear and rotating flows with uniformly distributed PIs, whereas Experiments 3 and 4 feature bubble columns and sprays, which exhibit non-uniform PI distributions. By analyzing all four cases, it becomes possible to assess the validity of the uniform distribution assumption across a range of flow conditions.

For each experiment, the images were manually labeled by drawing bounding boxes (BBs) around each PI. The circular representation of each PI was then derived from the center coordinates and the width and height of the corresponding BB. In the case of Exp. 2, the images exhibit slight astigmatism, and the PIs were approximated as ellipses with principal axes defined by the BB dimensions. This mild astigmatism in Exp. 2 also provides an opportunity to assess the influence of weak optical distortion on the validity of the results." [SK25a]

6.3 Definition of Overlap Metrics

"To gain insights into the nature of PI overlap from the test dataset, robust measures, independent of used experimental parameters (i.e. non-dimensional), must be defined first for the evaluation." [SK25a]

6.3.1 Metrics for the Number of Overlaps

"To define the number of overlaps a PI j has with other PIs ℓ in an image, PI overlap must first be detected. To do this, the intersection-over-union (IoU) measure can be used, compare Fig. 6.4. The IoU for the two PIs

$$\text{IoU}(\text{PI}_\ell, \text{PI}_j) = \frac{A_{\text{PI},j} \cap A_{\text{PI},\ell}}{A_{\text{PI},j} \cup A_{\text{PI},\ell}} \quad (6.1)$$

can be computed directly from the two circular areas A_{PI} . An overlap between PI j and PI ℓ is then detected if the condition $\text{IoU}(\text{PI}_\ell, \text{PI}_j) > 0$ is met. This check is performed for every PI pairing in the image to construct an IoU matrix. The number of overlaps of PI_j with other PIs is given by the cardinality of the set

$$N_{\text{OL},j} = |\{\text{PI}_\ell \mid \text{IoU}(\text{PI}_\ell, \text{PI}_j) > 0\}| \quad (6.2)$$

of PIs that have an $\text{IoU} > 0$ with PI_j . The number of overlaps per PI $N_{\text{OL},j}$ indicates the degree of PI overlap that a given PI j experiences. The degree of the PI overlap, characterizing a cluster of overlapping PI is, therefore, described by $N_{\text{OL},j}$, with a first degree overlap describing the overlap of two PIs, a second degree corresponds to a triplet overlap cluster and so on. Finally, the fraction of PIs without any overlap, relative to the total number of PIs, N_{noOL}/N_P , is also of interest. The number of PIs that do not overlap with any other PIs is defined by the set

$$N_{\text{noOL}} = |\{\text{PI}_j \mid \text{IoU}(\text{PI}_\ell, \text{PI}_j) = 0\}|. \quad (6.3)$$

"[SK25a]

6.3.2 Metrics for the Amount of Overlaps

Metrics for First Degree Overlaps

"There are many measures to quantify the amount of PI overlap. Dreisbach *et al.* [48] used the IoU to characterize the amount of overlap between two PIs. The IoU is a good measure to describe a first degree overlap, when the overlap in relation to the whole first degree cluster is of interest. However, the IoU has the disadvantage of referencing the intersection area (i.e. the overlapped region) to the union area of the two PIs. This means that for the same intersection area, different IoU values can result depending on the size of the other PI, as illustrated in Fig. 6.4. This renders the IoU a non-robust measure when the focus is on the covered area of a specific PI. Moreover, the IoU does not provide direct information about the amount of area overlapped on PI j ; it only reflects the union area, which also varies with the distance between the PI centres.




Figure 6.4: "Different measures to describe first degree PI overlaps: IoU (left), IoM (middle) and IoA (right). The issues with the IoU and the IoM are visualized." [SK25a]. Modified figure from [SK25a].

An alternative measure is the Intersection over Minimum (IoM):

$$\text{IoM}(\text{PI}_\ell, \text{PI}_j) = \frac{A_{\text{PI},j} \cap A_{\text{PI},\ell}}{\min(A_{\text{PI},j}, A_{\text{PI},\ell})} \quad (6.4)$$

also known as the Szymkiewicz–Simpson Coefficient [204]. This measure was used by Sax *et al.* [SDLK23] to characterize PI overlap. The IoM has the advantage over the IoU of referencing the overlapped area to the area of only one PI, making it more focused on the extent of coverage. However, when two overlapping PIs j and ℓ are considered, the reference area changes depending on which PI is smaller, rather than consistently referring to PI j , which renders the measure problematic for the use in higher degree overlaps. This effect is illustrated in Fig. 6.4. To address this issue, the IoM is slightly modified. Instead of referencing the minimum area when evaluating the overlap of PI j with other PIs, the area $A_{\text{PI},j}$ of PI j is always used as the reference" [SK25a]. This leads to the IoA, previously introduced in Eq. (4.27) in Chapter 4. For improved readability and completeness, the definition is restated here.

$$\text{IoA}(\text{PI}_\ell, \text{PI}_j) = \frac{A_{\text{PI},j} \cap A_{\text{PI},\ell}}{A_{\text{PI},j}} \quad (6.5)$$

"The IoA has the advantage of consistently referring to the same PI in an PI overlap. The IoA is equivalent to the IoM for uniformly sized PIs. The intersection area of two circles of radii $r_{\text{PI}j}$ and $r_{\text{PI}\ell}$ (i.e. $d_{\text{PI}}/2$) with the Euclidean distance $d_{C,j\ell}$ between the circle centers can be derived from geometrical considerations [211]:

$$A_{\text{PI},j} \cap A_{\text{PI},\ell} = \begin{cases} r_{\text{PI}j}^2 \arccos\left(\frac{d_{C,j\ell}^2 + r_{\text{PI}j}^2 - r_{\text{PI}\ell}^2}{2d_{C,j\ell}r_{\text{PI}j}}\right) + r_{\text{PI}\ell}^2 \arccos\left(\frac{d_{C,j\ell}^2 + r_{\text{PI}\ell}^2 - r_{\text{PI}j}^2}{2d_{C,j\ell}r_{\text{PI}\ell}}\right) \\ \quad - \frac{1}{2} \sqrt{(-d_{C,j\ell} + r_{\text{PI}j} + r_{\text{PI}\ell})(d_{C,j\ell} + r_{\text{PI}j} - r_{\text{PI}\ell})} \\ \quad \sqrt{(d_{C,j\ell} - r_{\text{PI}j} + r_{\text{PI}\ell})(d_{C,j\ell} + r_{\text{PI}j} + r_{\text{PI}\ell})}, & \text{if } r_{\text{PI}j} + r_{\text{PI}\ell} > d_{C,j\ell} \\ 0, & \text{if } r_{\text{PI}j} + r_{\text{PI}\ell} \leq d_{C,j\ell} \end{cases} \quad (6.6)$$

More recently, Xu *et al.* [218] introduced an alternative measure to characterize PI overlap. They defined the overlap ratio (OLR) of a PI j as:

$$\text{OLR} = \frac{r_{\text{PI},j} + r_{\text{PI},\ell} - d_{C,j\ell}}{2 \min(r_{\text{PI},j}, r_{\text{PI},\ell})}. \quad (6.7)$$

This measure quantifies the difference between the touching distance of the two PIs, given by $r_{\text{PI},j} + r_{\text{PI},\ell}$, and their actual centre-to-centre distance, normalized by the diameter of the smaller PI. The OLR (with reference to the smaller PI) and the IoA are consistent in the sense that, for different combinations of PI sizes and distances yielding the same IoA, the OLR also remains constant. However, the two measures are not equivalent, i.e., $\text{IoA} \neq \text{OLR}$. While the OLR emphasizes the spatial separation between PIs, the IoA directly quantifies the overlapped area. A limitation of the OLR, similar to that of the IoM, is that it always normalizes with respect to the smaller PI, and thus does not consistently refer to the same PI across comparisons. For these reasons, the IoA is adopted in the present work." [SK25a]

Metrics for Higher Degree Overlaps

"The IoA can be used to characterize first degree PI overlaps, i.e. the overlap between two PIs. Higher degree overlaps are visualized in Fig. 6.5. To investigate such higher degree PI overlaps, i.e. between PI j and other PIs $\ell = 1, 2, \dots, N_{\text{OL}}$, further metrics must be defined. The $\text{IoA}_{\text{max}} =$

$\max(\text{IoA}(\text{PI}_1, \text{PI}_j), \dots, \text{IoA}(\text{PI}_{N_{\text{OL}}}, \text{PI}_j))$ provides a measure for the largest first degree PI overlap, that the PI j experiences in a higher degree cluster, compare middle part of Fig. 6.5.




Figure 6.5: "First and higher degree PI overlaps visualized (left). Metrics for higher degree PI overlaps are visualized (right). The FRA and TCA are shown on the example of a PI in IPI."[SK25a]. Modified figure from [SK25a].

For IPI, the total fraction of the PI that is covered, or its complement, the remaining free area, is often of interest, as it provides insight into how much of the interference pattern can be used for undisturbed evaluation (see the right example of Fig. 6.5).

However, simply summing every IoA of PI j with any other PI $\ell = 1, 2, \dots, N_{\text{OL}}$ overestimates the covered PI area as areas repeatedly overlapped are also counted multiple times, compare Fig. 6.6. To determine the free remaining area (FRA), multiple times overlapped areas of the PI should only be considered once.




Figure 6.6: "A higher degree PI overlap visualized. The gray area is the PI for which the PI overlap is regarded. The red area is the intersection area, and the blue area shows the twice overlapped area of the PI. The summed IoA overestimated the TCA of the PI by the blue area."[SK25a]. Modified figure from [SK25a].

A robust measure for this case is the total covered area (TCA) and its complement the FRA of PI j . The total covered area is defined by the union area of all first degree intersection areas that PI j has with other PIs ℓ . Instead of using the absolute value of the covered area, the fraction of the covered area over the area of PI j is used

$$\text{TCA}_j = \frac{\bigcup_{\ell \neq j} (A_{\text{PI},j} \cap A_{\text{PI},\ell})}{A_{\text{PI},j}} \quad (6.8)$$

as it places the $\text{TCA}_j \in [0, 1]$. The complement FRA can then be simply defined as

$$\text{FRA}_j = 1 - \text{TCA}_j \quad (6.9)$$

the remaining area fraction. The TCA and FRA are visualized in Fig. 6.5.

To compute the TCA, the union of the intersection areas must be known. This can be computed from the IoA in principle using the inclusion-exclusion principle of Poincaré and Sylvester [74]. However, the computation requires the intersection areas between all combinations of PIs from the first to the highest degree overlap. Since this is computationally unfeasible, a quadrature-like approach is used to compute the TCA in the following."[SK25a]

6.3.3 Direct Computation of Overlapped Particle Image Fraction

"The TCA is computed using a quadrature-like approach, in which the circle area is discretized into pixels on a Cartesian grid (see Fig. 6.7). The circle areas are approximated by pixel counting, analogous to a Riemann sum. The grid consists of pixels $(u, v) \in \mathbb{Z}^2$. In this method, each PI is defined by its in-plane position $(x_{\text{PI}}, y_{\text{PI}})$ and diameter d_{PI} , and is represented as a binary circle mask

$$\mathcal{M}_{\text{circle}}(u, v) = \begin{cases} 1, & \text{if } (u - x_{\text{PI}})^2 + (v - y_{\text{PI}})^2 \leq (\frac{d_{\text{PI}}}{2})^2 \\ 0, & \text{otherwise} \end{cases} \quad (6.10)$$

where the mask takes the value 1 for pixels inside the circle and 0 elsewhere. By summing these binary masks over a zero-initialized image, a combined mask is formed in which each pixel value indicates the number of overlapping circles. The TCA for a given PI j is approximated by counting the number of pixels within its own circle mask that have a value greater than one, and dividing this by the total number of pixels in that mask.




Figure 6.7: "Deterministic pixel counting approach by simulating discretized circles. A circle consists of pixels with a value of one (marked as gray). The overlap areas of each circle are obtained from the number of pixels with a value of larger than one (marked as red). Different grid resolutions of the simulation are shown from 100 px to 400 px."[SK25a]. Figure from [SK25a].

The set of pixels belonging to PI j is defined as:

$$\Omega_{\text{PI}} = \left\{ (u, v) \in \mathbb{Z}^2 \mid (u - x_{\text{PI}})^2 + (v - y_{\text{PI}})^2 \leq (\frac{d_{\text{PI}}}{2})^2 \right\}. \quad (6.11)$$

The TCA is then approximated by

$$\text{TCA}_j \approx \frac{\sum_{u, v \in \Omega_{\text{PI}, j}} \mathbf{1}(\mathcal{M}_{\text{comb}}(u, v) > 1)}{|\Omega_{\text{PI}, j}|} \quad (6.12)$$

where $\mathbf{1}$ is the indicator function, equal to 1 if the condition is true and 0 otherwise, and

$$\mathcal{M}_{\text{comb}}(u, v) = \sum_{\ell=1}^{N_{\text{P}}} \mathcal{M}_{\text{circle}, \ell}(u, v) \quad (6.13)$$

is the combined mask of all N_P PIs.

This circle mask method enables direct approximation of the TCA and FRA. However, since the circle areas are discretized onto a pixel grid, the method inherently introduces a discretization error. As it follows a quadrature-like approach, increasing the grid resolution improves the accuracy of the approximation, and the error asymptotically vanishes with finer discretization. To ensure that the chosen grid resolution was sufficiently fine, the method was executed four times using different grid sizes $(u, v) \in \{1, \dots, N_{\text{px}}\} \times \{1, \dots, N_{\text{px}}\}$, with $N_{\text{px}} = \{100, 200, 300, 400\}$; see Fig. 6.7. Figure 6.8 shows the TCA computed for $N_P = 60$ PIs as a function of the summed PI area relative to the image area. A slight difference is observed between the 200 px and 300 px resolutions, while almost no difference is observed between 300 px and 400 px; see Fig. 6.8. This indicates that a resolution of $N_{\text{px}} = 400$ is likely sufficient to obtain a reliable approximation of the TCA."[SK25a]




Figure 6.8: "The approximation of the TCA by the quadrature-like method for $N_P = 60$. The discretized simulations are shown for different grid resolutions $N_{\text{px}} = \{200, 300, 400\}$. It can be seen that almost no difference between the $N_{\text{px}} = 300$ and $N_{\text{px}} = 400$ resolution can be observed."[SK25a]. Figure from [SK25a].

6.3.4 Definition of a Scaling Metric

"Before evaluating the overlap metrics as a function of the number of PIs and PI size, a scaling metric that consolidates both N_P and z_P into a single value is introduced. Cierpka *et al.* [34] introduced a modified definition of seeding density \mathcal{S} , as a ratio of the total PI area A_{PI} to the total image area A_{image} . This definition

$$\mathcal{S} = \sum_{\ell=1}^{N_P} \frac{A_{\text{PI},\ell}}{A_{\text{image}}} \quad (6.14)$$

is also employed in the present work to summarize both the number and size of PIs. The seeding density defined in Eq. (6.14) relates the PI area to the image area, which is determined by the region of interest (RoI) of the respective experiment. Since the RoI is generally unknown, it is here defined as the smallest enclosing rectangle (reflecting the typical rectangular shape of camera

sensors) that contains all PIs ℓ across all images, i.e., across all time steps t . The reference area is thus defined as:

$$A_{\text{image}} := \left[\max \left(x_{\text{PI},\ell,t} + \frac{d_{\text{PI},\ell,t}}{2} \right) - \min \left(x_{\text{PI},\ell,t} - \frac{d_{\text{PI},\ell,t}}{2} \right) \right] \left[\max \left(y_{\text{PI},\ell,t} + \frac{d_{\text{PI},\ell,t}}{2} \right) - \min \left(y_{\text{PI},\ell,t} - \frac{d_{\text{PI},\ell,t}}{2} \right) \right] \quad (6.15)$$

The seeding density \mathcal{S} depends

$$\mathcal{S} \sim N_{\text{P}} d_{\text{PI}}^2 \sim N_{\text{P}} z_{\text{P}}^2 \quad (6.16)$$

linearly on the number of PIs, N_{P} , and quadratically on the PI diameter, d_{PI} . Given the linear relationship between the distance of the PI to the focal plane, z_{P} , and d_{PI} under sufficient defocusing [57, 146], \mathcal{S} is therefore also quadratically dependent on z_{P} . The same quadratic dependence applies to the aperture diameter, i.e., $\mathcal{S} \sim D_{\text{a}}^2$.

Consequently, \mathcal{S} substitutes the two dimensional parameter space spanned by N_{P} and z_{P} to a one-dimensional parameter that is independent of physical units. From Eq. (6.16), it follows that various combinations of N_{P} and z_{P} can yield the same seeding density. For instance, an increased number of PIs in the image requires a reduction in defocus by the square root to maintain the same seeding density. Conversely, increased defocus must be compensated by a quadratic reduction in the number of PIs. For an arbitrary scaling $c \in \mathbb{R}$, this relationship between N_{P} and z_{P} can be expressed as $N_{\text{P}} z_{\text{P}}^2 = c N_{\text{P}} (z_{\text{P}}/\sqrt{c})^2$. By computing the defined metrics over \mathcal{S} , the experimentalist can select an appropriate combination of N_{P} and z_{P} for a given seeding density."[SK25a]

6.4 Scaling Model for Particle Image Overlap

"The means of all previously introduced measures, $N_{\text{OL},j}$, N_{noOL} , IoA_{max} , TCA, and FRA, are computed for the 10×6 parameter space of the test dataset. This way the influence of PI size and number of PIs on the PI overlap can be investigated. Additionally, are the measures computed for the validation image sets.

An exemplary plot of the number of overlaps per PI, N_{OL} , as a function of N_{P} and z_{P} is shown in Figs. 6.9a and 6.9b, respectively. It can be seen that the number of overlaps scales linearly with N_{P} and quadratically with z_{P} . However, the proportional scaling of N_{OL} with respect to N_{P} varies depending on the z -position, and vice versa. While this allows for an assessment of the general behavior of this overlap measure in relation to the experimental parameters, it does not provide a universal scaling law. In contrast, Fig. 6.10a presents the number of overlaps per PI plotted against the seeding density \mathcal{S} , which combines both parameters N_{P} and z_{P} into a single metric. It becomes immediately apparent that, when scaled by \mathcal{S} , the individual curves collapse onto a single master curve. As illustrated in Figs. 6.10 and 6.11, this behavior is consistent across all previously defined measures. This observation suggests that \mathcal{S} acts as a robust scaling parameter for PI overlap, indicating a scaling law that is independent of the specific measurement setup.

The second observation that can be drawn is that, the four experiments show a good alignment with the theoretical model, as the data points follow the metric curves closely, compare Figs. 6.10 and 6.11. This shows that despite all four image sets stemming from experiments with different flow topology, the assumption of uniform distribution describes the distribution of PIs in the image sufficiently well. It also shows that despite using equally sized PIs, the results can be extended to the PI of various sizes from the four experiments.




Figure 6.9: "The mean number of overlaps per PI over the number of PIs for different defocus distances (a) and over the defocus distance for different number of PIs (b)."[SK25a]. Figure from [SK25a].

The third observation that can be drawn is that the scaling laws also capture the PI overlap behavior observed in Exp. 3, which exhibits a small amount of astigmatism. It can therefore be assumed that the results also hold for weak astigmatism (aspect ratio < 1.66) in the PIs. However, no definitive statement for stronger astigmatism can be made. "[SK25a]

Number Overlaps

"Fig. 6.10a shows the mean number of overlaps per PI j (i.e. the degree of PI overlap) in an image as a function of the seeding density, \mathcal{S} . It can be observed that the expected number of overlaps per PI, $N_{OL,j}$, increases linearly with \mathcal{S} . The trend follows a straight line with a slope of four, indicating that at a critical seeding density of $\mathcal{S} = 0.25$, an average of one overlap per PI can be expected. In other words: the average PI cluster contains two PIs, or the average PI overlap is of degree one.




Figure 6.10: "The mean number of overlaps per PI over the seeding density $\mathcal{S} \sim N_P z_P^2$ for different number of PIs in the image (a). The mean fraction of PIs without overlap in an image over the seeding density \mathcal{S} and for different number of PIs in the image (b). Shown are also the fits to the curves (dashed line --), which are specified in Tab.6.1. The values obtained from the validation experiments are provided by the markers ($\triangle, \triangleleft, \diamond, \triangleright$). Images from the validation experiments are shown in Fig. 6.3."[SK25a]. Figure from [SK25a].

Fig. 6.10b shows the mean fraction of PIs in the image that experience no overlap. This measure is particularly interesting, as it provides insight into the proportion of PIs that can be processed more reliably. The number of overlap-free PIs N_{noOL} decreases rapidly for small values of \mathcal{S} , with the slope gradually leveling off. This is a significant finding, as it demonstrates that even slight defocusing near the focal plane has a substantial impact on PI overlap, whereas increasing the seeding density beyond $\mathcal{S} = 0.75$ has minimal effect on the number of overlap-free PIs. Additionally, Fig. 6.10b shows that the decline of the fraction of overlap-free PIs follows a Pareto function."[SK25a]

First Degree Overlaps Characteristics

"Most detection algorithms are capable of distinguishing PIs involved in multiple small first-degree overlaps, but typically struggle with larger first-degree overlaps. The total overlapped PI area within a cluster does not indicate whether the area is covered by several small or a few large first-degree overlaps. Most studies focusing on PI detection define overlap metrics exclusively for first-degree PI overlaps [48, SDLK23, 218, 158]. Therefore, when PI edges are of interest, such as in PI detection for DPTV, it is useful to consider first-degree PI overlap metrics in conjunction with higher-degree metrics. The maximum IoA is a particularly insightful metric, as it reflects the strongest first-degree overlap within a cluster. When combined with the TCA, it provides insight into how individual PI overlaps contribute to the total covered area, i.e., whether a PI is primarily covered by a single large overlap or by multiple smaller ones.

Fig. 6.11a shows the change of the mean IoA_{max} per PI with \mathcal{S} . For low seeding densities a steep incline of the max IoA can be observed, which gradually becomes less steep with increasing \mathcal{S} . It can be seen that the increase in max IoA is approximately linear for $\mathcal{S} < 0.25$ with a slope of approx one."[SK25a]




Figure 6.11: "The max IoA per PI over the seeding density $\mathcal{S} \sim N_P z_P^2$ for different number of PIs in the image (a). The mean TCA with $N_{\text{px}} = 400$ over the seeding density \mathcal{S} for different number of PIs in the image (b). Shown is also the fit to the curves (dashed line ---), which is specified in Tab.6.1. The values obtained from the validation experiments are provided by the markers (\triangle , \triangleleft , \diamond , \triangleright). Images from the validation experiments are shown in Fig. 6.3."[SK25a]. Figure from [SK25a].

Free Particle Area in Higher Degree PI Overlaps

"The extent of overlap in higher-degree PI interactions is described by the TCA and FRA, and constitutes a particularly important measure for IPI, where the PI area is used rather than the edges. Fig. 6.11b shows how the TCA scales with the seeding density \mathcal{S} . The TCA increases steeply at low seeding densities and asymptotically approaches one. Most practical applications will not involve seeding densities greater than $\mathcal{S} = 1$, and thus will not encounter mean TCAs exceeding approximately 70%. In fact, in many cases, PI overlap should be minimized as much as possible. Therefore, if an average TCA of e.g. 10% is considered acceptable, the seeding density should be kept below $\mathcal{S} = 0.09$."[SK25a]

The goal of this chapter "is to provide practical guidance for experimenters on selecting the number of particles and defocus lengths, to ensure that specific limits on PI overlap are maintained. The fit functions for all evaluated measures—obtained via least squares regression—are presented in Tab.6.1. These functions enable the creation of simple plots for each measure, so that experimenters can use the PI overlap model without the need to repeat the empirical investigation."[SK25a]

Table 6.1: "The different metrics to describe overlap and their respective models are shown. The model type as well as the fit function and the parameters are given. The goodness of fit is measured by R-squared."[SK25a].
Table from [SK25a].

Metric	Model	Model-function	Parameters	R-squared
$\text{mean}(N_{\text{noOL}}/N_{\text{P}})$	Pareto (PDF)	$\text{mean}(N_{\text{noOL}}/N_{\text{P}}) = \frac{1}{c_1} (1 + c_2 \frac{(\mathcal{S}-c_3)}{c_1})^{(-\frac{1}{c_2}-1)}$	$c_1 = 0.193$ $c_2 = 0.156$ $c_3 = -0.307$	99.999%
$\text{mean}(N_{\text{OL}})$	linear	$\text{mean}(N_{\text{OL}}) = c_1 \mathcal{S}$	$c_1 = 4.063$	99.840%
$\text{mean}(\text{IoA}_{\text{max}})$	rational function	$\text{mean}(\text{IoA}_{\text{max}}) = \frac{c_1 \mathcal{S}}{1+c_2 \mathcal{S}}$	$c_1 = 1.245$ $c_2 = 1.460$	99.999%
$\text{mean}(\text{TCA})$	rational function	$\text{mean}(\text{TCA}) = \frac{c_1 \mathcal{S}}{1+\mathcal{S}}$	$c_1 = 1.324$	99.956%

6.5 Concluding Remarks on Particle Image Overlap

"PI overlap presents a major challenge in DPTV, where incomplete PI boundary information increases detection miss rates and uncertainty in z -position estimation. In IPI, it, furthermore, reduces the usable PI area for fringe pattern analysis, which is essential for accurately determining the particle diameter (d_{P}). The developed model, which describes PI overlap independently of the optical setup, particle number, or PI size, is based on the assumption of a uniform spatial distribution of PIs with uniform sizes. Despite these simplifying assumptions, a comparison with experimental data, featuring non-uniform particle distributions and multiple PI sizes, shows that the model captures the observed behavior with sufficient accuracy. The model also performs well under conditions of mild astigmatism (i.e. aspect ratios of 1.66 or smaller). However, its validity under stronger astigmatic distortions remains uncertain and may require further investigation. The empirical study of PI overlap demonstrates that the seeding density $\mathcal{S} \sim N_{\text{P}} z_{\text{P}}^2$ serves as a powerful scaling parameter for describing PI overlap independently of the optical system. It was further observed that the PI overlap remains constant for any combination of N_{P} and z_{P} that results in a constant seeding density. This implies that, for example, the number of particles can be increased

linearly while the amount of defocusing is reduced by the square root, without altering the mean PI overlap.

The number of overlaps experienced by a PI, scales linearly with the seeding density, with a critical threshold at $\mathcal{S} = 0.25$, where an average of one overlap per PI can be expected. This means that if the PIs cover 25% of the recorded image, the average overlap cluster degree is one. Small variations in defocus length near the focal plane have a significant impact on the fraction of overlap-free PIs, particularly because $\mathcal{S} \sim z_p^2$. At approximately $\mathcal{S} \approx 0.16$, only about half of the PIs can be expected to remain overlap-free. As the seeding density increases further, its effect on the number of overlap-free PIs begins to saturate. Finally, the scaling of the TCA and FRA can be approximated using a quadrature-like method. The relation of the TCA and FRA to the seeding density S , provides crucial information on the expected fraction of a PI area which is not under influence of overlap and can be used in IPI, without restrictions. Knowledge of expected values for key metrics, such as the average degree of an overlap cluster, the number of overlap-free PIs and the TCA, enables experimenters to design their setups (i.e. number of particles, defocus length, and aperture diameter) in a way that maintains PI overlap within acceptable limits on average. The derived scaling laws for each overlap metric serve as a look up table to design experiments. With the established scaling laws, critical seeding densities can be identified. These critical values enable the selection of appropriate combinations of PI sizes and particle numbers tailored to the specific requirements of a given experiment. This approach helps to ensure that experimental data with appropriate PI overlap can be acquired in a quantitative manner instead of relying on an experimenters estimation from experience."[SK25a]

7 Conclusions and Outlook

Gas-liquid DTPFs are central to numerous natural and industrial processes, including weather systems, combustion, and chemical reactors. The accurate measurement of these flows is essential for advancing theoretical understanding and improving safety, efficiency, and sustainability. Yet, capturing such gas-liquid DTPFs, especially in three dimensions and under optically constrained conditions, remains a significant challenge. Conventional optical techniques often fall short due to limitations in spatial resolution, simultaneous tracking of both phases, and accurate particle sizing. Alternative methods, such as MRI, may offer solutions but typically require specialised equipment that is not always readily accessible.

This thesis explored how the DPTV and IPI techniques can be utilised and extended to enable simultaneous 3D position determination of tracers and dispersed phase particles, along with accurate sizing of the dispersed phase.

DPTV and IPI were selected as the foundational techniques, as they require only standard PIV imaging setups. DPTV enables 3D position determination with only a single optical access, while IPI offers high sizing accuracy for particles in large fields of view. First, IPI was extended to the backscatter regime using the DSE of the LMT, and the theoretical model was experimentally validated. Second, DPTV and IPI were combined to achieve simultaneous 3D position and size determination of dispersed bubbles or droplets. This was realised through an IP approach, which was validated on both synthetic and experimental data, and benchmarked against existing methods. Third, DPTV was adapted for two-phase measurements without additional equipment by employing CNNs to distinguish between tracers and bubbles or droplets. This approach was validated across five experiments and one synthetic dataset. Finally, the challenge of PI overlap was addressed. Although the IP approach is robust to overlap, it requires a reliable detection algorithm for initialisation. A statistical model was therefore developed to estimate the extent of PI overlap under varying experimental conditions, helping to identify the operational limits beyond which overlap may compromise data quality. The model and its assumptions were validated on four separate experiments.

Summary of Key Findings and Reflection on the Research Questions

The central research objective was divided into five smaller research questions, each of which is addressed in a dedicated chapter of this thesis. As each chapter concludes with its own findings, only the main insights are summarised here. The overarching conclusions are discussed in the next section.

RQ1 examined how IPI can be extended to the backscatter regime for single optical access. Using the DSE of the LMT and introducing the generalised visibility criterion, suitable scattering angles can be identified for both bubbles and droplets. A transfer function for IPI at any scattering angle can be derived from finding the dominant scattering orders (visibility criterion) and using the glare point maps to determine the glare point spacing. For bubbles the $p = [0, 2]$ pairing and for droplet first the $p = [0, 3]$ and further in backscatter the $p = [0, 2]$ pairing dominate in the backscatter regime. It was shown that the polarisation has a significant impact on the visibility of fringes. A

validation experiment confirmed the feasibility of backscatter IPI. The method can be used for any real valued refractive index.

RQ2 addressed how uncertainties and limitations shift in side- and backscatter compared to front-scatter. Variations in the scattering angle and the resulting changes in glare point spacing affect fringe frequency, thereby altering the measurable size range and associated uncertainty. The ability to control this size range via the scattering angle is a key insight with implications beyond backscatter measurements. The fringe intensity also varies with the scattering angle, influencing the SNR. Small changes in the scattering angle can cause intensity fluctuations, introducing an in-plane position dependence in measurement uncertainty, which typically not occurs in front-scatter. The generally lower SNR in backscatter necessitates more advanced post-processing techniques and a carefully considered experimental design. The selection of an appropriate scattering angle is particularly challenging in the case of droplets, as it requires accounting for visibility, intensity, polarisation, and additional phenomena such as rainbow angles. In contrast, bubbles involve less complex optical phenomena, and thus the choice of scattering angle demands comparatively less meticulous consideration.

RQ3 explored how an inverse problem approach improves particle position and size estimation. Notably, it decouples the diameter from the depth estimation, enabling accurate sizing in 3D volumes. The method combines IPI and DPTV, thereby allowing for simultaneous 3D position and size determination in a unified approach. The IP approach performs at least as good as current benchmarks on real data, and results on synthetic data suggest a considerable increase in both position and size estimation accuracy. The method has also shown robustness against aliasing and noise. Another major improvement is its resistance to PI overlap, a key limitation in conventional IPI and DPTV methods. The approach can also be applied to position determination alone for tracers by omitting the size parameter. Its primary limitations include a high computational cost, which currently precludes real-time applications, and a strong dependence on near-complete initialisation for optimal performance. Additionally, deviations in the aperture shape can introduce inaccuracies in depth estimation. For full images containing multiple particles, the approach also showed reduced accuracy under poor initialisation. However, even poor initial guesses were still improved by the method.

RQ4 investigated whether CNNs can reliably distinguish between tracer particles and bubbles or droplets in two-phase DPTV for DTPFs. The proposed method uses a pattern based distinction of defocused PIs. The two networks tested achieved high accuracy (95-100%) in phase distinction, which remained consistent generalising across different datasets. While maximum recall varied more significantly (67-100%), precision remained high, resulting in few or no ghost particles. These results demonstrate that CNNs can reliably distinguish phases, enabling two-phase DPTV using post-processing alone, without the need for additional equipment. Although phase distinction using CNNs is highly reliable, wavelength-based separation may offer even greater accuracy but requires extra equipment. The CNN-based method also demands more computational resources. Notably, the networks exhibited a false positive class bias towards the dispersed phase, indicating that training data and loss functions must be carefully adjusted.

A method was also introduced to generate experiment-specific training data from unlabelled images, reducing the need for manual labelling. This enables the creation of larger, more suitable datasets for high-capacity models, in the context of DPTV and IPI. However, the approach requires a larger initial setup effort and expertise in machine learning.

RQ5 focused on quantifying expected PI overlap in DPTV and IPI based on experimental parameters. Robust measures were developed for both first- and higher-degree overlaps, with the seeding

density \mathcal{S} (defined as the ratio of total PI area to image area) confirmed as a reliable scaling parameter. A statistical model was introduced to estimate key overlap characteristics, including the fraction of non-overlapping PIs, the expected degree of overlap clusters, and the non-overlapped PI area. The derived scaling laws are non-dimensional and thus applicable to any experimental setup. This model helps define operational limits for reliable data acquisition. While the model holds for polydisperse particle sizes and weak astigmatism, its applicability to strong astigmatism remains uncertain.

With the main research objective divided into five subsidiary research questions, Fig. 7.1 illustrates how these questions are interconnected. Chapter 3 introduced a method to render IPI in backscatter configurations viable, akin to DPTV. This enables the combined use of DPTV and IPI for simultaneous three-dimensional position and size determination. However, due to the inherently lower SNR in backscatter, IPI requires more advanced evaluation techniques. This challenge is addressed in Chapter 4, where the developed IP approach demonstrates strong resistance to noise, even when applied to experimental images presented in Chapter 3.




Figure 7.1: Overview of the measurement framework for DTPFs in optically constrained environments. A CNN is used for the initial detection and phase separation. The BBs outputted by the CNN are used to initialise the IP approach. Once the IP approach converges, the final particle positions and sizes are derived from the updated input vector. The contributions of each chapter are indicated in grey within the schematic. Modified figure adapted from [SGK25].

When combining IPI and DPTV in backscatter, several issues must be resolved: the sizing of three-dimensionally distributed particles, simultaneous position determination, and PI overlap. The IP approach offers improvements in all these areas. The primary limitation of backscatter measurements lies in the low SNR. The IP method mitigates this constraint, as its resistance to aliasing permits measurements closer to the focal plane, thereby allowing to significantly enhance the SNR of the raw data. Nevertheless, the IP approach requires initialisation, which can be provided by CNNs, as employed in Chapter 5 and in previous studies [48, SDLK23, 223, 159]. Although the CNNs in Chapter 5 did not achieve the recall rates necessary for the near-complete initialisation demanded by the IP method, other studies have demonstrated that such high recall values are

attainable [48, SDLK23], particularly when CNNs are trained for a specific experimental setup. However, obtaining training data tailored to a specific experiment is often challenging, typically requiring extensive manual labelling or synthetic data that may lack sufficient specificity. Chapter 5 addresses this issue by introducing the auto-labelling method. The IP approach described in Chapter 4 can be extended to two types of particles: tracers and bubbles or droplets. To enable this, the input vector must be expanded, as shown in Fig. 7.1. This expanded input vector necessitates not only position initialisation but also prior classification, which is addressed by the two-phase DPTV approach using CNNs in Chapter 5. Finally, CNNs used to initialise the IP approach operate effectively only within certain ranges of PI overlap [14, 48, SDLK23]. To ensure these limits are not exceeded, the scaling laws for PI overlap derived in Chapter 6 can be applied.

Practical Implications and Broader Impact

Combined, this thesis provides the foundational building blocks necessary to enable simultaneous 3D position determination of tracers and dispersed phase particles, along with sizing of the dispersed phase, in optically restricted environments. Deriving these metrics concurrently allows for the measurement of instationary flows and yields more holistic data compared to separate measurements. The ability to obtain such data despite optical limitations facilitates in-situ measurements in real industrial applications. Moreover, it enables the extraction of information, such as slip velocities, that would otherwise be difficult to access. As these techniques build upon standard PIV imaging equipment, they are broadly applicable in most laboratories and do not require specialised or costly apparatus, thus making advanced flow investigations more accessible on a limited budget.

While the combined approach opens new avenues for broader applications, the insights from individual chapters also have wider relevance. The finding that the measurable size range and uncertainties in IPI can be actively controlled by adjusting the scattering angle broadens the scope of potential applications. This enables the measurement of smaller particles and the tuning of uncertainties based on physical constraints - particularly useful for studying satellite droplets in breakup processes or minute bubbles near nucleation. The IP approach permits higher seeding densities in both IPI and DPTV, enhancing the spatial resolution of velocity vectors in DPTV. This qualification for higher seeding densities addresses one of the main drawbacks of the DPTV method. Increased seeding density in IPI enables the study of denser particle swarms, previously inaccessible to IPI. This is particularly valuable for investigating sprays or dense bubble swarms. Performing IPI in 3D also allows for more comprehensive size distribution measurements compared to planar techniques, which do not capture out-of-plane phenomena. While the IP approach can increase the seeding density in IPI, this is only feasible to a certain extent. As the seeding density rises, multi-particle scattering emerges as a new limitation, one not previously encountered, as PI overlap would have prevented seeding densities susceptible to this effect. Finally, the development of scaling laws for PI overlap provides practical guidance for any defocusing-based technique. This approach bridges the gap between algorithmic limitations and experimental design, enabling seeding densities to be increased up to known algorithmic thresholds without reliance on guesswork. It facilitates the targeted use of higher source densities, thereby maximising the achievable spatial resolution of velocity vectors in DPTV. It also helps determine the maximum usable PI sizes, which enhances the accuracy of depth position estimation by enabling the identification of optimal defocus sensitivities, and it improves the reliability of particle size evaluation in IPI. Overall, this enables the potential of the evaluation algorithms to be more fully utilised.

Outlook and Direction of Future Work

Building upon the methods developed and insights derived in this thesis, future research can progress in several directions. A logical and technically feasible next step involves the integration of CNNs for two-phase DPTV with the IP approach. This integration would enable the realisation of the conceptual framework proposed in this work. The combination is straightforward, as the BBs generated by the CNNs can be directly utilised as input for the IP method. The only requirement is that the input vector distinguishes between tracers (x_P, y_P, z_P) and dispersed phase particles such as bubbles or droplets (x_P, y_P, z_P, d_P), as illustrated in Fig. 7.1. This combined approach could be validated analogously to the methodology presented in Chapter 4. However, it is essential to ensure that the intensity of tracers is approximately equal to that of the dispersed phase particles to avoid detection bias. Furthermore, the optimisation vector within the IP framework may be skewed towards brighter particles. To mitigate this, careful selection of scattering angles and tracer types is required, guided by intensity plots as demonstrated in Chapter 3.

Given that backscatter IPI is predominantly constrained by low SNR, future research should prioritise the development of noise-resilient PI detection and evaluation techniques. Although the IP approach exhibits robustness to noise, it remains dependent on accurate PI detection. The most straightforward approach to improve the PI detection and evaluation is improved background subtraction. One potential solution involves the application of neural networks for image denoising, which could be employed to generate initial region proposals for PI detection. Such networks may be trained end-to-end on unlabelled data by synthetically introducing various types of noise (e.g., Gaussian noise, salt-and-pepper noise) into raw experimental images. Promising architectures for this task include denoising auto encoders [205], CNNs, and GANs. Furthermore, enhancing image quality prior to the PI evaluation should include background subtraction using pre-recorded background images, as well as the application of adaptive filtering techniques.

The scaling laws for PI overlap derived in this thesis are currently valid only for low degrees of astigmatism. Extending these models to accommodate stronger astigmatism represents a valuable direction for future work. Within the existing framework, astigmatism can be incorporated by simulating ellipses rather than circles. Due to the optical characteristics of cylindrical lenses, only the ratio of the principal axes, rather than their orientation, would constitute an additional degree of freedom. This theoretical extension is relatively straightforward, and computational cost is not a limiting factor at a resolution of $N_{px} = 400$. Nevertheless, the increased complexity would necessitate a significantly larger volume of labelled experimental data for validation. Successfully extending the scaling laws to include astigmatism would substantially broaden their applicability, particularly in the context of APTV measurements.

Further development of two-phase DPTV using CNNs also presents a promising research direction. Recent studies [14, 48, SDLK23, 159, 223] have demonstrated the efficacy of CNNs in PI detection and their capacity to achieve high recall rates. The auto-labelling approach introduced in Chapter 5 could be leveraged to generate larger, experiment-specific datasets. This would facilitate the use of more complex CNN architectures, potentially enhancing both detection and classification performance. To support this advancement, it may be beneficial to develop and disseminate the auto-labelling code, thereby enabling broader adoption.

Another advancement would be for authors to share their labelled experimental data, enabling the DPTV and APTV community to build large and diverse training sets. Such datasets would further advance PI detection but would require a coordinated effort. While this would involve an initial investment of effort, it would benefit the community in the long term, as the acquisition of training data remains one of the primary hurdles limiting the broader use of CNNs in this field.

In summary, this thesis lays the foundation for a framework enabling in-situ diagnostics of DTPFs in environments where conventional measurement techniques are inadequate. Several novel concepts have been introduced to address key trade-offs typically encountered in complex measurement scenarios. These developments support the acquisition of more accurate data and open access to measurements that were previously unattainable due to equipment limitations. The methodologies presented herein facilitate more detailed experimental investigations of DTPFs, thereby contributing to a deeper theoretical understanding of the underlying physical processes and supporting the advancement of more efficient and sustainable technologies.

8 References

Bibliography

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from <https://www.tensorflow.org/>.
- [2] M. Abramowitz and I. A. Stegun. *Handbook of mathematical functions*. Dover, New York N.Y., 1965. ISBN 978-0486612720.
- [3] R. J. Adrian. Scattering particle characteristics and their effect on pulsed laser measurements of fluid flow: speckle velocimetry vs particle image velocimetry. *Applied Optics*, 23(11):1690–1691, Jun 1984. DOI 10.1364/AO.23.001690. URL <https://doi.org/10.1364/AO.23.001690>.
- [4] R. J. Adrian and J. Westerweel. *Particle Image Velocimetry*. Cambridge University Press, Cambridge, United Kingdom, 2010. ISBN 9780521440080. URL <https://www.cambridge.org/9780521440080>.
- [5] H.-E. Albrecht, M. Borys, N. Damaschke, and C. Tropea. *Laser Doppler and Phase Doppler Measurement Techniques*. Springer Berlin, Heidelberg, Germany, 2003. DOI 10.1007/978-3-662-05165-8. URL <https://doi.org/10.1007/978-3-662-05165-8>.
- [6] M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017. URL <https://arxiv.org/abs/1701.04862>.
- [7] T. Atherton and D. Kerbyson. Using phase to represent radius in the coherent circle Hough transform. In *IEE Colloquium on Hough Transforms*, pages 5/1–5/4, May 1993. URL <https://ieeexplore.ieee.org/document/243198>.
- [8] T. Atherton and D. Kerbyson. Size invariant circle detection. *Image and Vision Computing*, 17(11):795–803, Sep 1999. ISSN 0262-8856. DOI 10.1016/S0262-8856(98)00160-7. URL [https://doi.org/10.1016/S0262-8856\(98\)00160-7](https://doi.org/10.1016/S0262-8856(98)00160-7).
- [9] T. J. Atherton and D. J. Kerbyson. The coherent circle Hough transform. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 27.1–27.10, 1993. DOI 10.5244/C.7.27. URL <http://dx.doi.org/10.5244/C.7.27>.
- [10] O. Baker. Simultaneous flow of oil and gas. *Oil and Gas Journal*, 53(12):185–195, 1954.
- [11] S. Balachandar. *Fundamentals of Dispersed Multiphase Flows*. Cambridge University Press, Cambridge, 2024. ISBN 9781009160469. DOI 10.1017/9781009160452. URL <https://doi.org/10.1017/9781009160452>.

- [12] R. Barnkob and M. Rossi. General defocusing particle tracking: fundamentals and uncertainty assessment. *Experiments in Fluids*, 61, 2020. DOI 10.1007/s00348-020-2937-5. URL <https://doi.org/10.1007/s00348-020-2937-5>.
- [13] R. Barnkob, C. J. Kähler, and M. Rossi. General defocusing particle tracking. *Lab on a chip*, 15(17):3556–3560, 2015. DOI 10.1039/c5lc00562k. URL <http://dx.doi.org/10.1039/c5lc00562k>.
- [14] R. Barnkob, C. Cierpka, M. Chen, S. Sachs, P. Mäder, and M. Rossi. Defocus particle tracking: a comparison of methods based on model functions, cross-correlation, and neural networks. *Measurement Science and Technology*, 32(9):094011, Jun 2021. ISSN 0957-0233. DOI 10.1088/1361-6501/abfef6. URL <http://dx.doi.org/10.1088/1361-6501/abfef6>.
- [15] P. Baues. Huygens’ principle in inhomogeneous, isotropic media and a general integral equation applicable to optical resonators. *Opto-Electronics*, 1:37–44, Feb 1969. DOI 10.1007/BF01476791. URL <https://doi.org/10.1007/BF01476791>.
- [16] E. H. Bocanegra, N. Dam, D. van der Voort, G. Bertens, and W. van de Water. Measuring droplet size distributions from overlapping interferometric particle images. *Review of Scientific Instruments*, 86(2):023709, Feb 2015. ISSN 0034-6748. DOI 10.1063/1.4909537. URL <https://doi.org/10.1063/1.4909537>.
- [17] A. Bochkovskiy, C.-Y. Wang, and H.-y. Liao. Yolov4: Optimal speed and accuracy of object detection, Apr 2020. URL <https://doi.org/10.48550/arXiv.2004.10934>.
- [18] C. F. Bohren and D. R. Huffman. *Absorption and Scattering of Light by Small Particles*. WILEY-VCH, Weinheim, Germany, 1st edition, 1998. ISBN 9783527618156. DOI 10.1002/9783527618156. URL <https://doi.org/10.1002/9783527618156>. Originally published in 1983.
- [19] M. Born and E. Wolf. *Principles of Optics: 60th Anniversary Edition*. Cambridge University Press, 2019. ISBN 9781108477437. DOI 10.1017/9781108769914. URL <https://doi.org/10.1017/9781108769914>.
- [20] M. Brandestini. Topoflow - a digital full range Doppler velocity meter. *IEEE Transactions on Sonics and Ultrasonics*, 25(5):287–292, Sep 1978. DOI 10.1109/T-SU.1978.31029. URL <https://doi.org/10.1109/T-SU.1978.31029>.
- [21] C. E. Brennen. *Fundamentals of Multiphase Flow*. Cambridge University Press, 2005. ISBN 0521 848040. DOI 10.1017/CBO9780511807169. URL <https://doi.org/10.1017/CBO9780511807169>.
- [22] W. Brouwer. Matrix methods in optical instrument design. In *Lecture notes and supplements in physics*, New York, 1964. W.A. Benjamin.
- [23] M. Brunel, S. Coetmellec, G. Gréhan, and H. Shen. Interferometric out-of-focus imaging simulator for irregular rough particles. *Journal of the European Optical Society -Rapid Publication*, 9, 2014. DOI 10.2971/jeos.2014.14008. URL <http://dx.doi.org/10.2971/jeos.2014.14008>.
- [24] P. J. Bryanston-Cross, M. Funes-Gallanzi, C. Quan, and T. R. Judge. Holographic particle image velocimetry (hpiv). *Optics & Laser Technology*, 24(5):251–256, 1992. ISSN 00303992. DOI 10.1016/0030-3992(92)90066-B. URL [http://dx.doi.org/10.1016/0030-3992\(92\)90066-B](http://dx.doi.org/10.1016/0030-3992(92)90066-B).
- [25] K. Bryden, K. W. Ragland, and S.-C. Kong. *Combustion Engineering*. CRC Press, Boca Raton, 3 edition, 2022. DOI 10.1201/b22232. URL <https://doi.org/10.1201/b22232>.

- [26] J. Butler and R. Bonnecaze. Imaging of particle shear migration with electrical impedance tomography. *Physics of Fluids - PHYS FLUIDS*, 11:1982–1994, Aug 1999. DOI 10.1063/1.870062. URL <https://doi.org/10.1063/1.870062>.
- [27] R. H. Byrd, J. C. Gilbert, and J. Nocedal. A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming*, 89:149–185, 2000. DOI 10.1007/PL00011391. URL <https://doi.org/10.1007/PL00011391>.
- [28] S. L. Ceccio and D. L. George. A review of electrical impedance techniques for the measurement of multiphase flows. *Journal of Fluids Engineering*, 118(2):391–399, Jun 1996. ISSN 0098-2202. DOI 10.1115/1.2817391. URL <https://doi.org/10.1115/1.2817391>.
- [29] R. Cerqueira, E. Paladino, B. Ynumaru, and C. Maliska. Image processing techniques for the measurement of two-phase bubbly pipe flows using particle image and tracking velocimetry (PIV/PTV). *Chemical Engineering Science*, 189:1–23, 2018. ISSN 0009-2509. DOI 10.1016/j.ces.2018.05.029. URL <https://doi.org/10.1016/j.ces.2018.05.029>.
- [30] Y. Cha, W. Choi, G. Suh, S. Mahmoudkhani, and O. Buyukozturk. Autonomous structural visual inspection using region-based deep learning for detecting multiple damage types. *Computer-Aided Civil and Infrastructure Engineering*, 33:731–747, Nov 2017. DOI 10.1111/mice.12334. URL <https://doi.org/10.1111/mice.12334>.
- [31] R. Chen and L.-S. Fan. Particle image velocimetry for characterizing the flow structure in three-dimensional gas-liquid-solid fluidized beds. *Chemical Engineering Science*, 47(13):3615–3622, 1992. ISSN 0009-2509. DOI 10.1016/0009-2509(92)85077-O. URL [https://doi.org/10.1016/0009-2509\(92\)85077-O](https://doi.org/10.1016/0009-2509(92)85077-O).
- [32] C. Cierpka and C. J. Kähler. Particle imaging techniques for volumetric three-component (3d3c) velocity measurements in microfluidics. *Journal of Visualization*, 15(1):1–31, Feb 2012. ISSN 1343-8875. DOI 10.1007/s12650-011-0107-9. URL <https://doi.org/10.1007/s12650-011-0107-9>.
- [33] C. Cierpka, R. Segura, R. Hain, and C. J. Kähler. A simple single camera 3c3d velocity measurement technique without errors due to depth of correlation and spatial averaging for microfluidics. *Measurement Science and Technology*, 21(4):045401, Mar 2010. DOI 10.1088/0957-0233/21/4/045401. URL <https://dx.doi.org/10.1088/0957-0233/21/4/045401>.
- [34] C. Cierpka, M. Rossi, R. Segura, and C. J. Kähler. On the calibration of astigmatism particle tracking velocimetry for microflows. *Measurement Science and Technology*, 22(1):015401, 2011. ISSN 0957-0233. DOI 10.1088/0957-0233/22/1/015401. URL <http://dx.doi.org/10.1088/0957-0233/22/1/015401>.
- [35] C. Cierpka, J. König, M. Chen, D. Boho, and P. Mäder. On the use of machine learning algorithms for the calibration of astigmatism ptv. In *Proceedings of the 13th International Symposium on Particle Image Velocimetry (PIV19)*, 2019.
- [36] S. A. Collins. Lens-system diffraction integral written in terms of matrix optics. *Journal of the Optical Society of America*, 60(9):1168–1177, 1970. DOI 10.1364/JOSA.60.001168. URL <https://doi.org/10.1364/JOSA.60.001168>.
- [37] Y. Cui, Y. Song, C. Sun, A. Howard, and S. Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4109–4118, 2018. DOI 10.1109/CVPR.2018.00432. URL <http://dx.doi.org/10.1109/CVPR.2018.00432>.

- [38] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2:303–314, 1989. DOI <https://doi.org/10.1007/BF02551274>. URL <https://doi.org/10.1007/BF02551274>.
- [39] J. Dainty, J. W. Goodman, G. Parry, T. S. McKechnie, M. Françon, and A. E. Ennos. *Laser Speckle and Related Phenomena*. Springer Nature, Topics in Applied Physics, 1975. DOI 10.1007/978-3-662-43205-1. URL <https://doi.org/10.1007/978-3-662-43205-1>.
- [40] P. Debye. Das elektromagnetische Feld um einen Zylinder und die Theorie des Regenbogens. *Physikalische Zeitschrift*, 9:1072–1079, 1908.
- [41] R. Deepak Selvakumar and S. Dhinakaran. Forced convective heat transfer of nanofluids around a circular bluff body with the effects of slip velocity using a multi-phase mixture model. *International Journal of Heat and Mass Transfer*, 106:816–828, 2017. ISSN 0017-9310. DOI 10.1016/j.ijheatmasstransfer.2016.09.108. URL <https://doi.org/10.1016/j.ijheatmasstransfer.2016.09.108>.
- [42] S. Dehaeck and J. van Beeck. Multifrequency interferometric particle imaging for gas bubble sizing. *Experiments in Fluids*, 45:823–831, Apr 2008. DOI 10.1007/S00348-008-0502-8. URL <https://doi.org/10.1007/S00348-008-0502-8>.
- [43] E. Delnoij, J. Kuipers, W. van Swaaij, and J. Westerweel. Measurement of gas–liquid two-phase flow in bubble columns using ensemble correlation piv. *Chemical Engineering Science*, 55(17):3385–3395, 2000. ISSN 0009-2509. DOI 10.1016/S0009-2509(99)00595-3. URL [https://doi.org/10.1016/S0009-2509\(99\)00595-3](https://doi.org/10.1016/S0009-2509(99)00595-3).
- [44] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. DOI 10.1109/CVPR.2009.5206848. URL <https://doi.org/10.1109/CVPR.2009.5206848>.
- [45] S. Deotale, S. Dutta, J. Moses, V. Balasubramaniam, and A. Chinnaswamy. Foaming characteristics of beverages and its relevance to food processing. *Food Engineering Reviews*, 12: 229–250, Jun 2020. DOI 10.1007/s12393-020-09213-4. URL <https://doi.org/10.1007/s12393-020-09213-4>.
- [46] S. Discetti and Y. Liu. Machine learning for flow field measurements: a perspective. *Measurement Science and Technology*, 34(2):021001, Nov 2022. DOI <https://dx.doi.org/10.1088/1361-6501/ac9991>. URL 10.1088/1361-6501/ac9991.
- [47] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, Oct 2020. URL <https://doi.org/10.48550/arXiv.2010.11929>.
- [48] M. Dreisbach, R. Leister, M. Probst, P. Friederich, A. Stroh, and J. 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, Sep 2022. ISSN 0957-0233. DOI 10.1088/1361-6501/ac8a09. URL <http://dx.doi.org/10.1088/1361-6501/ac8a09>.
- [49] V. Dumoulin and F. Visin. A guide to convolution arithmetic for deep learning, Mar 2018. URL <https://doi.org/10.48550/arXiv.1603.07285>.

-
- [50] C. Dunker, C. Roloff, and A. Grassmann. Interferometric laser imaging for in-flight cloud droplet sizing. *Measurement Science and Technology*, 27(12), Oct 2016. DOI 10.1088/0957-0233/27/12/124004. URL <https://dx.doi.org/10.1088/0957-0233/27/12/124004>.
 - [51] E. Ebert, A. Kleinwächter, R. Kostbade, and N. A. Damaschke. Interferometric particle imaging for cavitation nuclei characterization in cavitation tunnels and in the wake flow. In *Proceedings of the 17th International Symposium on Application of Laser Techniques to Fluid Mechanics*, Lisbon, Portugal, Jul 2014.
 - [52] G. E. Elsinga, F. Scarano, B. Wieneke, and B. W. van Oudheusden. Tomographic particle image velocimetry. *Exp Fluids*, 41:933–947, 2006. DOI 10.1007/s00348-006-0212-z. URL <https://doi.org/10.1007/s00348-006-0212-z>.
 - [53] M. A. Erinin, B. Néel, M. T. Mazzatenta, J. H. Duncan, and L. Deike. Comparison between shadow imaging and in-line holography for measuring droplet size distributions. *Experiments in Fluids*, 64(96), 2023. DOI s00348-023-03633-8. URL <https://doi.org/10.1007/s00348-023-03633-8>.
 - [54] E. O. Etuke and R. T. Bonnecaze. Measurement of angular velocities using electrical impedance tomography. *Flow Measurement and Instrumentation*, 9(3):159–169, 1998. ISSN 0955-5986. DOI 10.1016/S0955-5986(98)00020-X. URL [https://doi.org/10.1016/S0955-5986\(98\)00020-X](https://doi.org/10.1016/S0955-5986(98)00020-X).
 - [55] A. Fiaz, K. Aftab, B. Qiu, M. Sultan, and J. Ma. Advancements of spraying technology in agriculture. In F. Ahmad and M. Sultan, editors, *Technology in Agriculture*, chapter 2. IntechOpen, Rijeka, 2021. DOI 10.5772/intechopen.98500. URL <https://doi.org/10.5772/intechopen.98500>.
 - [56] S. Franchini and S. Krevor. Cut, overlap and locate: a deep learning approach for the 3d localization of particles in astigmatic optical setups. *Experiments in Fluids*, 61(6), Jun 2020. ISSN 0723-4864. DOI 10.1007/s00348-020-02968-w. URL <http://dx.doi.org/10.1007/s00348-020-02968-w>.
 - [57] T. Fuchs, R. Hain, and C. J. Kähler. In situ calibrated defocusing ptv for wall-bounded measurement volumes. *Measurement Science and Technology*, 27(8):084005, 2016. ISSN 0957-0233. DOI 10.1088/0957-0233/27/8/084005. URL <http://dx.doi.org/10.1088/0957-0233/27/8/084005>.
 - [58] N. Fujisawa, A. Hosokawa, and S. Tomimatsu. Simultaneous measurement of droplet size and velocity field by an interferometric imaging technique in spray combustion. *Measurement Science and Technology*, 14:1341–1349, Jul 2003. DOI 10.1088/0957-0233/14/8/320. URL <https://doi.org/10.1088/0957-0233/14/8/320>.
 - [59] D. A. Gabor. New microscopic principle. *Nature*, 161:777–778, 1948. DOI 10.1038/161777a0. URL <https://doi.org/10.1038/161777a0>.
 - [60] A. Gerrard and J. M. Burch. *Introduction to Matrix Methods in Optics*. A Wiley-Interscience Publication. John Wiley & Sons, London and New York, 1975. ISBN 9780471296850.
 - [61] J. Gire, L. Denis, C. Fournier, É. Thiébaud, F. Soulez, and C. Ducottet. Digital holography of particles: Benefits of the ‘inverse problem’ approach. *Measurement Science and Technology*, 19:074005, May 2008. DOI 10.1088/0957-0233/19/7/074005. URL <https://doi.org/10.1088/0957-0233/19/7/074005>.

- [62] R. Girshick. Fast R-CNN. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1440–1448, 2015. DOI 10.1109/ICCV.2015.169. URL <https://doi.org/10.1109/ICCV.2015.169>.
- [63] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 580–587, 2014. DOI 10.1109/CVPR.2014.81. URL <https://doi.org/10.1109/CVPR.2014.81>.
- [64] 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. *Appl. Opt.*, 34(36):8409–8421, Dec 1995. DOI 10.1364/AO.34.008409. URL <https://doi.org/10.1364/AO.34.008409>.
- [65] C. Gonzalez Viejo, D. Torrico, F. Dunshea, and S. Fuentes. Bubbles, foam formation, stability and consumer perception of carbonated drinks: A review of current, new and emerging technologies for rapid assessment and control. *Foods*, 8, Nov 2019. DOI 10.3390/foods8120596. URL <https://doi.org/10.3390/foods8120596>.
- [66] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [67] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks, 2014. URL <https://arxiv.org/abs/1406.2661>.
- [68] G. Gouesbet. Debye series formulation for generalized Lorenz-Mie theory with the Bromwich method. *Particle & Particle Systems Characterization*, 20(6):382–386, 2003. ISSN 09340866. DOI 10.1002/ppsc.200300886. URL <https://doi.org/10.1002/ppsc.200300886>.
- [69] 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, 2014. URL <https://arxiv.org/abs/1412.5567>.
- [70] Y. Hardalupas, S. Sahu, A. Taylor, and K. Zarogoulidis. Simultaneous planar measurement of droplet velocity and size with gas phase velocities in a spray by combined ilids and piv techniques. *Experiments in Fluids*, 49:417–434, Aug 2010. DOI 10.1007/s00348-009-0802-7. URL <https://doi.org/10.1007/s00348-009-0802-7>.
- [71] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1904–1916, 2015. DOI 10.1109/TPAMI.2015.2389824. URL <https://doi.org/10.1109/TPAMI.2015.2389824>.
- [72] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. DOI 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- [73] O. S. Heavens, R. W. Ditchburn, and S. George. *Insight into Optics, Chapter 4.13*. American Journal of Physics, 1 edition, 1992. URL <https://doi.org/10.1119/1.17089>. URL: <https://doi.org/10.1119/1.17089>.
- [74] N. Henze. *Die Formel des Ein- und Ausschließens*, pages 72–77. Vieweg+Teubner Verlag, Wiesbaden, 2011. ISBN 978-3-8348-8649-1. DOI 10.1007/978-3-8348-8649-1_11. URL https://doi.org/10.1007/978-3-8348-8649-1_11.

-
- [75] Y. D. Herlambang, Kurnianingsih, A. Roihatin, T. Prasetyo, Marliyati, Taufik, and J.-C. Shyu. A numerical study of bubble blockage in microfluidic fuel cells. *Processes*, 10(5), 2022. ISSN 2227-9717. DOI 10.3390/pr10050922. URL <https://doi.org/10.3390/pr10050922>.
 - [76] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS 2020)*, 2020. URL <https://arxiv.org/abs/2006.11239>.
 - [77] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9:1735–1780, Nov 1997. DOI 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
 - [78] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989. ISSN 0893-6080. DOI 10.1016/0893-6080(89)90020-8. URL [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).
 - [79] P. V. C. Hough. Machine analysis of bubble chamber pictures. In *Proceedings, 2nd International Conference on High-Energy Accelerators and Instrumentation, HEACC 1959*, volume C590914, pages 554–558, CERN, Geneva, Switzerland, Sep 1959. URL <https://s3.cern.ch/inspire-prod-files-5/53d80b0393096ba4afe34f5b65152090>.
 - [80] E. A. Hovenac and J. A. Lock. Assessing the contributions of surface waves and complex rays to far-field Mie scattering by use of the Debye series. *Journal of the Optical Society of America A*, 9(5):781–795, 1992. ISSN 1084-7529. DOI 10.1364/JOSAA.9.000781. URL <https://doi.org/10.1364/JOSAA.9.000781>.
 - [81] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In F. Bach and D. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR.
 - [82] M. L. Jakobsen, W. J. Easson, C. A. Greated, and D. H. Glass. Particle image velocimetry: simultaneous two-phase flow measurements. *Measurement Science and Technology*, 7(9):1270, Sep 1996. DOI 10.1088/0957-0233/7/9/012. URL <https://dx.doi.org/10.1088/0957-0233/7/9/012>.
 - [83] N. Jegham, C. Y. Koh, M. Abdelatti, and A. Hendawi. Yolo evolution: A comprehensive benchmark and architectural review of yolov12, yolov11, and their previous versions, 2025. URL <https://doi.org/10.48550/arXiv.2411.00201>.
 - [84] P. Jenneson, R. Luggar, E. Morton, O. Gundogdu, and U. Tuzun. Examining nanoparticle assemblies using high spatial resolution x-ray microtomography. *Journal of Applied Physics*, 96:2889–2894, Oct 2004. DOI 10.1063/1.1776635. URL <https://doi.org/10.1063/1.1776635>.
 - [85] D. D. Joseph and D. Ocando. Slip velocity and lift. *Journal of Fluid Mechanics*, 454:263–286, 2002. DOI 10.1017/S0022112001007145. URL <https://doi.org/10.1017/S0022112001007145>.
 - [86] H. Kao and A. Verkman. Tracking of single fluorescent particles in three dimensions: use of cylindrical optics to encode particle position. *Biophysical Journal*, 67(3):1291–1300, 1994. ISSN 0006-3495. DOI 10.1016/S0006-3495(94)80601-0. URL [https://doi.org/10.1016/S0006-3495\(94\)80601-0](https://doi.org/10.1016/S0006-3495(94)80601-0).
 - [87] H. P. Kao and A. S. Verkman. Tracking of single fluorescent particles in three dimensions: use of cylindrical optics to encode particle position. *Biophysical Journal*, 67(3):1291–1300, 1994. ISSN 00063495. DOI 10.1016/S0006-3495(94)80601-0. URL [http://dx.doi.org/10.1016/S0006-3495\(94\)80601-0](http://dx.doi.org/10.1016/S0006-3495(94)80601-0).

- [88] N. Karmarkar. A new polynomial-time algorithm for linear programming-ii. *Combinatorica*, 4:373–395, Dec 1984. DOI 10.1007/BF02579150. URL <https://doi.org/10.1007/BF02579150>.
- [89] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2019. DOI 10.48550/arXiv.1812.04948. URL <https://doi.org/10.48550/arXiv.1812.04948>.
- [90] T. Kawaguchi, Y. Akasaka, and M. Maeda. Size measurements of droplets and bubbles by advanced interferometric laser imaging technique. *Measurement Science and Technology*, 13(3):308, Feb 2002. DOI 10.1088/0957-0233/13/3/312. URL <https://dx.doi.org/10.1088/0957-0233/13/3/312>.
- [91] D. Khalitov and E. Longmire. Simultaneous two-phase piv by two-parameter phase discrimination. *Experiments in Fluids*, 32:252–268, Feb 2002. DOI 10.1007/s003480100356. URL <https://doi.org/10.1007/s003480100356>.
- [92] V. Khare and H. M. Nussenzveig. Theory of the rainbow. *Physical Review Letters*, 33:976–980, 1974. DOI 10.1103/PHYSREVLETT.33.976. URL <https://doi.org/10.1103/PHYSREVLETT.33.976>.
- [93] J. J. Kiejar, Y. Wu, S. Coëtmellec, D. Lebrun, G. Gréhan, and M. Brunel. Size determination of mixed liquid and frozen water droplets using interferometric out-of-focus imaging. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 178:108–116, 2016. ISSN 0022-4073. DOI 10.1016/j.jqsrt.2015.09.009. URL <https://doi.org/10.1016/j.jqsrt.2015.09.009>.
- [94] K. T. Kiger and C. Pan. Piv technique for the simultaneous measurement of dilute two-phase flows. *Journal of Fluids Engineering*, 122(4):811–818, Jul 2000. ISSN 0098-2202. DOI 10.1115/1.1314864. URL <https://doi.org/10.1115/1.1314864>.
- [95] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations (ICLR)*, 2015. DOI 10.48550/arXiv.1412.6980. URL <https://doi.org/10.48550/arXiv.1412.6980>.
- [96] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. In *2nd International Conference on Learning Representations (ICLR)*, 2014. URL <https://arxiv.org/abs/1312.6114>.
- [97] A. Kitagawa, Y. Hagiwara, and T. Kouda. Ptv investigation of phase interaction in dispersed liquid-liquid two-phase turbulent swirling flow. *Experiments in Fluids*, 42:871–880, Jun 2007. DOI 10.1007/s00348-007-0291-5. URL <https://doi.org/10.1007/s00348-007-0291-5>.
- [98] 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. ISSN 00218502. DOI 10.1016/0021-8502(86)90063-7. URL [https://doi.org/10.1016/0021-8502\(86\)90063-7](https://doi.org/10.1016/0021-8502(86)90063-7).
- [99] J. König, M. Chen, W. Rösing, D. Boho, P. Mäder, and C. Cierpka. On the use of a cascaded convolutional neural network for three-dimensional flow measurements using astigmatic ptv. *Measurement Science and Technology*, 31(7):074015, May 2020. ISSN 0957-0233. DOI 10.1088/1361-6501/ab7bfd. URL <http://dx.doi.org/10.1088/1361-6501/ab7bfd>.
- [100] W. Kosiwczuk, A. Cessou, M. Trinité, and B. Lecordier. Simultaneous velocity field measurements in two-phase flows for turbulent mixing of sprays by means of two-phase piv. *Experiments in Fluids*, 39:895–908, 2005. DOI 10.1007/s00348-005-0027-3. URL <https://doi.org/10.1007/s00348-005-0027-3>.

-
- [101] M. Kramer, H. Chanson, and S. F. and. Can we improve the non-intrusive characterization of high-velocity air–water flows? Application of lidar technology to stepped spillways. *Journal of Hydraulic Research*, 58(2):350–362, 2020. DOI 10.1080/00221686.2019.1581670. URL <https://doi.org/10.1080/00221686.2019.1581670>.
 - [102] T. M. Kreis, M. Adams, and W. P. O. Jueptner. Methods of digital holography: a comparison. In C. Gorecki, editor, *Optical Inspection and Micromasurements II*, volume 3098, pages 224–233. International Society for Optics and Photonics, SPIE, 1997. DOI 10.1117/12.281164. URL <https://doi.org/10.1117/12.281164>.
 - [103] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. *Neural Information Processing Systems*, 25:1097–1105, Jan 2012. DOI 10.1145/3065386. URL <https://doi.org/10.1145/3065386>.
 - [104] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, May 2017. ISSN 0001-0782. DOI 10.1145/3065386. URL <https://doi.org/10.1145/3065386>.
 - [105] C. Kähler, S. Scharnowski, and C. Cierpka. On the uncertainty of digital piv and ptv near walls. *Experiments in Fluids*, 52:1641–1656, 2012. DOI 10.1007/s00348-012-1307-3. URL <https://doi.org/10.1007/s00348-012-1307-3>.
 - [106] G. Lacagnina, S. Grizzi, M. Falchi, F. D. Felice, and G. P. Romano. Simultaneous size and velocity measurements of cavitating microbubbles using interferometric laser imaging. *Experiments in Fluids*, 50:1153–1167, Feb 2011. DOI 10.1007/S00348-011-1055-9. URL <https://doi.org/10.1007/S00348-011-1055-9>.
 - [107] P. Laven. Mieplot. <http://philiplaven.com/mieplot.htm>. URL <http://philiplaven.com/mieplot.htm>. "[Online; last accessed 15th Jun 2025]".
 - [108] Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and L. Jackel. Handwritten digit recognition with a back-propagation network. In D. Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 2, pages 396–404. Morgan-Kaufmann, 1989. URL https://proceedings.neurips.cc/paper_files/paper/1989/file/53c3bce66e43be4f209556518c2fcb54-Paper.pdf.
 - [109] Y. Lecun, B. Boser, J. Denker, D. Henderson, Howard R., H. Wayne, and L. Jackel. Handwritten digit recognition with a backpropagation network. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 2, 1990. URL <https://proceedings.neurips.cc/paper/1989/file/53c3bce66e43be4f209556518c2fcb54-Paper.pdf>.
 - [110] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. DOI 10.1109/5.726791. URL <https://doi.org/10.1109/5.726791>.
 - [111] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. ISSN 00189219. DOI 10.1109/5.726791. URL <http://dx.doi.org/10.1109/5.726791>.
 - [112] Y. LeCun, L. Bottou, G. B. Orr, and K. R. Müller. *Efficient BackProp*, pages 9–50. Springer Berlin Heidelberg, Berlin, Heidelberg, 1998. ISBN 978-3-540-49430-0. DOI 10.1007/3-540-49430-8_2. URL https://doi.org/10.1007/3-540-49430-8_2.
 - [113] R. Leister. *The Fluid Flow in an Open Wet Clutch*. PhD thesis, Karlsruher Institut für Technologie (KIT), 2023. URL <http://dx.doi.org/10.5445/IR/1000154812>.

- [114] R. Leister and J. Kriegseis. 3d-lif experiments in an open wet clutch by means of defocusing ptv. In *13th International Symposium on Particle Image Velocimetry (ISPIV 2019), Munich, Germany*, 2019. DOI 10.5445/IR/1000098119. URL <http://dx.doi.org/10.5445/IR/1000098119>.
- [115] R. Leister, T. Fuchs, P. Mattern, and J. Kriegseis. Flow-structure identification in a radially grooved open wet clutch by means of defocusing particle tracking velocimetry. *Experiments in Fluids*, 62(2), 2021. ISSN 0723-4864. DOI 10.1007/s00348-020-03116-0. URL <http://dx.doi.org/10.1007/s00348-020-03116-0>.
- [116] R. Leister, T. Fuchs, and J. Kriegseis. Defocusing ptv applied to an open wet clutch: from macro to micro. *Experiments in Fluids*, 64(94), 2023. DOI 10.1007/s00348-023-03623-w. URL <https://doi.org/10.1007/s00348-023-03623-w>.
- [117] W. Lentz. Generating Bessel functions in Mie scattering calculations using continued fractions. *Applied Optics*, 15:668–671, Mar 1976. DOI 10.1364/AO.15.000668. URL <https://doi.org/10.1364/AO.15.000668>.
- [118] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár. Microsoft coco: Common objects in context // microsoft coco: Common objects in context. In *European Conference on Computer Vision, 2014*, 2015. DOI 10.48550/arXiv.1405.0312. URL <http://dx.doi.org/10.48550/arXiv.1405.0312>.
- [119] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 936–944, 2017. DOI 10.1109/CVPR.2017.106. URL <http://dx.doi.org/10.1109/CVPR.2017.106>.
- [120] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar. Focal loss for dense object detection. In *IEEE Transactions on Pattern Analysis and Machine Intelligence 42*, volume 2017, pages 2999–3007, 2017. DOI 10.1109/ICCV.2017.324. URL <http://dx.doi.org/10.1109/ICCV.2017.324>.
- [121] A. Lipson, S. G. Lipson, and H. Lipson. *Optical Physics*. Cambridge University Press, 4 edition, 2010. DOI 10.1017/CBO9780511763120. URL <https://doi.org/10.1017/CBO9780511763120>.
- [122] S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia. Path aggregation network for instance segmentation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8759–8768, 2018. DOI 10.1109/CVPR.2018.00913. URL <https://doi.org/10.1109/CVPR.2018.00913>.
- [123] R. Luggar, E. Morton, P. Jenneson, and M. Key. X-ray tomographic imaging in industrial process control. *Radiation Physics and Chemistry*, 61(3):785–787, Jun 2001. ISSN 0969-806X. DOI S0969-806X(01)00404-2. URL [https://doi.org/10.1016/S0969-806X\(01\)00404-2](https://doi.org/10.1016/S0969-806X(01)00404-2).
- [124] A. I. Lvovsky. Fresnel equations. In *Encyclopedia of Optical Engineering*. Taylor and Francis, New York, 2013. DOI 10.1081/E-EOE-120047133. URL <http://dx.doi.org/10.1081/E-EOE-120047133>.
- [125] A. Maas, A. Hannun, and A. Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, 2013.

- [126] H. Maas, A. Gruen, and D. Papantoniou. Particle tracking velocimetry in three-dimensional flows. *Experiments in Fluids*, 15:133–146, 1993. DOI 10.1007/BF00190953. URL <https://doi.org/10.1007/BF00190953>.
- [127] M. Maeda, T. Kawaguchi, and K. Hishida. Novel interferometric measurement of size and velocity distributions of spherical particles in fluid flows. *Measurement Science and Technology*, 11, 2000. DOI 10.1088/0957-0233/11/12/101. URL <https://doi.org/10.1088/0957-0233/11/12/101>. Cited by: 141.
- [128] M. Maeda, Y. Akasaka, and T. Kawaguchi. Improvements of the interferometric technique for simultaneous measurement of droplet size and velocity vector field and its application to a transient spray. *Experiments in Fluids*, 33:125–134, Jan 2002. DOI 10.1007/s00348-002-0453-4. URL <https://doi.org/10.1007/s00348-002-0453-4>.
- [129] G. Martin, R. Castrejon-Pita, and I. Hutchings. Holographic measurement of drop-on-demand drops in flight. In *International Conference on Digital Printing Technologies*, volume 27, Oct 2011. DOI 10.2352/ISSN.2169-4451.2011.27.1.art00055_2. URL https://doi.org/10.2352/ISSN.2169-4451.2011.27.1.art00055_2.
- [130] MathWorks. Introduction to deep learning: what are convolutional neural networks?, 2017. URL <https://www.mathworks.com/videos/introduction-to-deep-learning-what-are-convolutional-neural-networks--1489512765771.html>. Last Accessed: June 16th 2025.
- [131] I. MathWorks. Matlab - gan. <https://www.mathworks.com/help/deeplearning/ug/train-generative-adversarial-network.html>, 2020. URL <https://www.mathworks.com/help/deeplearning/ug/train-generative-adversarial-network.html>. [Online; last accessed 15th Jun 2025].
- [132] K. Matsuura, K. Zarogoulidis, Y. Hardalupas, A. Taylor, T. Kawaguchi, D. Sugimoto, and K. Hishida. Simultaneous planar measurement of size and three-component velocity of droplets in an aero-engine airblast fuel spray by stereoscopic interferometric laser imaging technique. In *Proceedings of the 10th International Conference on Liquid Atomization and Spray Systems (ICLASS)*, 2006.
- [133] C. Matzler. Matlab functions for Mie scattering and absorption. Technical report, Institute of Applied Physics, University of Bern, Jun 2002. URL <https://omlc.org/software/mie/maetzlermie/Maetzler2002.pdf>.
- [134] F. D. McLeod and M. Anliker. A multiple gate pulsed directional Doppler flowmeter. In *IEEE Ultrasonics Symposium Proceedings*, Miami, FL, USA, 1971. IEEE.
- [135] M. Mehdizadeh Youshanlouei and M. Rossi. Deep learning and defocus imaging for determination of three-dimensional position and orientation of microscopic objects. *Physics of Fluids*, 36(8):082006, Aug 2024. ISSN 1070-6631. DOI 10.1063/5.0219081. URL <https://doi.org/10.1063/5.0219081>.
- [136] S. Meissner, J. Herold, L. Kirsten, C. Schneider, and E. Koch. 3d optical coherence tomography as new tool for microscopic investigations of nucleate boiling on heated surfaces. *International Journal of Heat and Mass Transfer*, 55(21):5565–5569, 2012. ISSN 0017-9310. DOI 10.1016/j.ijheatmasstransfer.2012.05.039. URL <https://doi.org/10.1016/j.ijheatmasstransfer.2012.05.039>.
- [137] G. Mie. Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen. *Annalen der Physik*, 330(3):377–445, 1908. DOI 10.1002/andp.19083300302. URL <https://doi.org/10.1002/andp.19083300302>.

- [138] L. Mondy, A. Graham, A. Majumdar, and L. Bryant. Techniques of measuring particle motions in concentrated suspensions. *International Journal of Multiphase Flow*, 12(3):497–502, 1986. ISSN 0301-9322. DOI 10.1016/0301-9322(86)90021-2. URL [https://doi.org/10.1016/0301-9322\(86\)90021-2](https://doi.org/10.1016/0301-9322(86)90021-2).
- [139] C. Mounaïm-Rousselle and O. Pajot. Droplet sizing by mie scattering interferometry in a spark ignition engine. *Particle & Particle Systems Characterization*, 16:160–168, 1999. DOI 10.1002/(SICI)1521-4117(199908)16:4<160::AID-PPSC160>3.0.CO;2-G. URL [https://doi.org/10.1002/\(SICI\)1521-4117\(199908\)16:4<160::AID-PPSC160>3.0.CO;2-G](https://doi.org/10.1002/(SICI)1521-4117(199908)16:4<160::AID-PPSC160>3.0.CO;2-G).
- [140] V. Nair and G. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning*, volume 27, pages 807–814, Jun 2010.
- [141] J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950. DOI 10.1073/pnas.36.1.48. URL <https://doi.org/10.1073/pnas.36.1.48>.
- [142] M. Nazarathy and J. Shamir. First-order optics—a canonical operator representation: lossless systems. *Journal of the Optical Society of America*, 72(3):356–364, Mar 1982. DOI 10.1364/JOSA.72.000356. URL <https://doi.org/10.1364/JOSA.72.000356>.
- [143] K. Nishino, N. Kasagi, and M. Hirata. Three-dimensional particle tracking velocimetry based on automated digital image processing. *Journal of Fluids Engineering*, 111(4):384–391, Dec 1989. ISSN 0098-2202. DOI 10.1115/1.3243657. URL <https://doi.org/10.1115/1.3243657>.
- [144] Y. Niwa, Y. Kamiya, T. Kawaguchi, and M. Maeda. Bubble sizing by interferometric laser imaging. In *10th International Symposium on Application of Laser Techniques to Fluid Mechanics*, 2000.
- [145] H. M. Nussenzveig. High-frequency scattering by a transparent sphere. i. direct reflection and transmission. *Journal of Mathematical Physics*, 10(1):82–124, Jan 1969. ISSN 0022-2488. DOI 10.1063/1.1664764. URL <https://doi.org/10.1063/1.1664764>.
- [146] M. G. Olsen and R. J. Adrian. Out-of-focus effects on particle image visibility and correlation in microscopic particle image velocimetry. *Experiments in Fluids*, 29(7):166–174, 2000. ISSN 0723-4864. DOI 10.1007/s003480070018. URL <http://dx.doi.org/10.1007/s003480070018>.
- [147] OpenAI, J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, R. Avila, I. Babuschkin, S. Balaji, V. Balcom, P. Baltescu, H. Bao, M. Bavarian, J. Belgum, I. Bello, J. Berdine, G. Bernadett-Shapiro, C. Berner, L. Bogdonoff, O. Boiko, M. Boyd, A.-L. Brakman, G. Brockman, T. Brooks, M. Brundage, K. Button, T. Cai, R. Campbell, A. Cann, B. Carey, C. Carlson, R. Carmichael, B. Chan, C. Chang, F. Chantzis, D. Chen, S. Chen, R. Chen, J. Chen, M. Chen, B. Chess, C. Cho, C. Chu, H. W. Chung, D. Cummings, J. Currier, Y. Dai, C. Decareaux, T. Degry, N. Deutsch, D. Deville, A. Dhar, D. Dohan, S. Dowling, S. Dunning, A. Ecoffet, A. Eleti, T. Eloundou, D. Farhi, L. Fedus, N. Felix, S. P. Fishman, J. Forte, I. Fulford, L. Gao, E. Georges, C. Gibson, V. Goel, T. Gogineni, G. Goh, R. Gontijo-Lopes, J. Gordon, M. Grafstein, S. Gray, R. Greene, J. Gross, S. S. Gu, Y. Guo, C. Hallacy, J. Han, J. Harris, Y. He, M. Heaton, J. Heidecke, C. Hesse, A. Hickey, W. Hickey, P. Hoeschele, B. Houghton, K. Hsu, S. Hu, X. Hu, J. Huizinga, S. Jain, S. Jain, J. Jang, A. Jiang, R. Jiang, H. Jin, D. Jin, S. Jomoto, B. Jonn, H. Jun, T. Kaftan, Łukasz Kaiser, A. Kamali, I. Kanitscheider, N. S. Keskar, T. Khan, L. Kilpatrick, J. W. Kim, C. Kim, Y. Kim, J. H. Kirchner,

- J. Kiros, M. Knight, D. Kokotajlo, Łukasz Kondraciuk, A. Kondrich, A. Konstantinidis, K. Koscic, G. Krueger, V. Kuo, M. Lampe, I. Lan, T. Lee, J. Leike, J. Leung, D. Levy, C. M. Li, R. Lim, M. Lin, S. Lin, M. Litwin, T. Lopez, R. Lowe, P. Lue, A. Makanju, K. Malfacini, S. Manning, T. Markov, Y. Markovski, B. Martin, K. Mayer, A. Mayne, B. McGrew, S. M. McKinney, C. McLeavey, P. McMillan, J. McNeil, D. Medina, A. Mehta, J. Menick, L. Metz, A. Mishchenko, P. Mishkin, V. Monaco, E. Morikawa, D. Mossing, T. Mu, M. Murati, O. Murk, D. Mély, A. Nair, R. Nakano, R. Nayak, A. Neelakantan, R. Ngo, H. Noh, L. Ouyang, C. O’Keefe, J. Pachocki, A. Paino, J. Palermo, A. Pantuliano, G. Parascandolo, J. Parish, E. Parparita, A. Passos, M. Pavlov, A. Peng, A. Perelman, F. de Avila Belbute Peres, M. Petrov, H. P. de Oliveira Pinto, Michael, Pokorny, M. Pokrass, V. H. Pong, T. Powell, A. Power, B. Power, E. Proehl, R. Puri, A. Radford, J. Rae, A. Ramesh, C. Raymond, F. Real, K. Rimbach, C. Ross, B. Rotsted, H. Roussez, N. Ryder, M. Saltarelli, T. Sanders, S. Santurkar, G. Sastry, H. Schmidt, D. Schnurr, J. Schulman, D. Selsam, K. Sheppard, T. Sherbakov, J. Shieh, S. Shoker, P. Shyam, S. Sidor, E. Sigler, M. Simens, J. Sitkin, K. Slama, I. Sohl, B. Sokolowsky, Y. Song, N. Staudacher, F. P. Such, N. Summers, I. Sutskever, J. Tang, N. Tezak, M. B. Thompson, P. Tillet, A. Tootoonchian, E. Tseng, P. Tuggle, N. Turley, J. Tworek, J. F. C. Uribe, A. Vallone, A. Vijayvergiya, C. Voss, C. Wainwright, J. J. Wang, A. Wang, B. Wang, J. Ward, J. Wei, C. Weinmann, A. Welihinda, P. Welinder, J. Weng, L. Weng, M. Wiethoff, D. Willner, C. Winter, S. Wolrich, H. Wong, L. Workman, S. Wu, J. Wu, M. Wu, K. Xiao, T. Xu, S. Yoo, K. Yu, Q. Yuan, W. Zaremba, R. Zellers, C. Zhang, M. Zhang, S. Zhao, T. Zheng, J. Zhuang, W. Zhuk, and B. Zoph. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>.
- [148] S. Paguada, L. Batina, I. Buhan, and I. Armendariz. Being patient and persistent: Optimizing an early stopping strategy for deep learning in profiled attacks, 2021. URL <https://doi.org/10.48550/arXiv.2111.14416>.
- [149] S. Pasch, H. Lange, R. Leister, and J. Kriegseis. Flow measurements above a dbd plasma actuator array by means of defocusing ptv. In *Proceedings of the International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics*, volume 21, Jul 2024. DOI 10.55037/lxaser.21st.47. URL <https://doi.org/10.55037/lxaser.21st.47>.
- [150] R. L. Powell. Experimental techniques for multiphase flows. *Physics of Fluids*, 20, 2008. DOI 10.1063/1.2911023. URL <https://doi.org/10.1063/1.2911023>.
- [151] D. Pöppe. Evaluating the applicability of flow measurement techniques for characterizing the microdroplet flow of the aerosol-on-demand printer. Master’s thesis, Karlsruhe Institute of Technology, 2023.
- [152] N. Qian. On the momentum term in gradient descent learning algorithms. *Neural Networks*, 12:145–151, 1999. ISSN 0893-6080. DOI 10.1016/S0893-6080(98)00116-6. URL [https://doi.org/10.1016/S0893-6080\(98\)00116-6](https://doi.org/10.1016/S0893-6080(98)00116-6).
- [153] L. Qieni, J. Wenhua, L. Tong, W. Xiang, and Z. Yimo. High-accuracy particle sizing by interferometric particle imaging. *Optics Communications*, 312:312–318, 2014. ISSN 0030-4018. DOI 10.1016/j.optcom.2013.09.049. URL <http://dx.doi.org/10.1016/j.optcom.2013.09.049>.
- [154] A. Quérel, P. Lemaitre, M. Brunel, E. Porcheron, and G. Gréhan. Real-time global interferometric laser imaging for the droplet sizing (ilids) algorithm for airborne research. *Measurement Science and Technology*, 21:015306, Dec 2009. DOI 10.1088/0957-0233/21/1/015306. URL <https://doi.org/10.1088/0957-0233/21/1/015306>.

- [155] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks, 2016. URL <https://doi.org/10.48550/arXiv.1511.06434>.
- [156] M. Raffel, C. E. Willert, F. Scarano, C. J. Kähler, S. T. Wereley, and J. Kompenhans. *Particle Image Velocimetry*. Springer International Publishing, Heidelberg, Germany, 2018. DOI 10.1007/978-3-319-68852-7. URL <https://doi.org/10.1007/978-3-319-68852-7>.
- [157] C. Raman. A new radiation. *Proceedings of the Indian Academy of Sciences - Section A*, 37(5):333–341, 1953. DOI 10.1007/BF03052651. URL <https://doi.org/10.1007/BF03052651>.
- [158] S. J. Rao, S. Sharma, S. Basu, and C. Tropea. Depth from defocus technique: a simple calibration-free approach for dispersion size measurement. *Experiments in Fluids*, 65, 2024. DOI 10.1007/s00348-024-03792-2. URL <https://doi.org/10.1007/s00348-024-03792-2>.
- [159] M. Ratz, S. Sachs, J. König, and C. Cierpka. A deep neural network architecture for reliable 3d position and size determination for lagrangian particle tracking using a single camera. *Measurement Science and Technology*, 34(10):105203, Jun 2023. DOI 10.1088/1361-6501/ace070. URL <https://dx.doi.org/10.1088/1361-6501/ace070>.
- [160] L. Rayleigh. On the light from the sky, its polarization and colour. *Philosophical Magazine*, 41(271):107–120, 1871.
- [161] J. Redmon and A. Farhadi. Yolo9000: Better, faster, stronger. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6517–6525, 2017. DOI 10.1109/CVPR.2017.690. URL <https://doi.org/10.1109/CVPR.2017.690>.
- [162] J. Redmon and A. Farhadi. Yolov3: An incremental improvement, 2018. URL <https://doi.org/10.48550/arXiv.1804.02767>.
- [163] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, 2016. DOI 10.1109/CVPR.2016.91. URL <https://doi.org/10.1109/CVPR.2016.91>.
- [164] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, Jun 2017. DOI 10.1109/TPAMI.2016.2577031. URL <http://dx.doi.org/10.1109/TPAMI.2016.2577031>.
- [165] M. Rossi. Synthetic image generator for defocusing and astigmatic piv/ptv. *Measurement Science and Technology*, 31(1):017003, 2020. ISSN 0957-0233. DOI 10.1088/1361-6501/ab42bb. URL <http://dx.doi.org/10.1088/1361-6501/ab42bb>.
- [166] M. Rossi and R. Barnkob. A fast and robust algorithm for general defocusing particle tracking. *Measurement Science and Technology*, 32(1):014001, 2020. ISSN 0957-0233. URL <http://dx.doi.org/10.1088/1361-6501/abad71>.
- [167] N. Roth, K. D. I. Anders, and A. Frohn. Determination of size, evaporation rate and freezing of water droplets using light scattering and radiation pressure. *Particle & Particle Systems Characterization*, 11:207–211, 1994. DOI 10.1002/ppsc.19940110307. URL <https://doi.org/10.1002/ppsc.19940110307>.
- [168] D. Rumelhart, G. Hinton, and R. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986. DOI <https://doi.org/10.1038/323533a0>. URL <https://doi.org/10.1038/323533a0>.

-
- [169] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115, Sep 2014. DOI 10.1007/s11263-015-0816-y. URL <https://doi.org/10.1007/s11263-015-0816-y>.
 - [170] P. Russell, J. A. Venning, B. W. Pearce, and P. A. Brandner. Calibration of Mie scattering imaging for microbubble measurement in hydrodynamic test facilities. *Experiments in Fluids*, 61, 2020. DOI 10.1007/s00348-020-2927-7. URL <https://doi.org/10.1007/s00348-020-2927-7>.
 - [171] D. K. Sarkar. *Thermal Power Plant*. Elsevier, 2015. DOI 10.1016/C2014-0-00536-9. URL <https://doi.org/10.1016/C2014-0-00536-9>.
 - [172] C. Sax. Particle detection by means of deep learning and conventional algorithms for defocusing particle tracking velocimetry data. Master’s thesis, Karlsruhe Institute of Technology (KIT), Institute of Fluid Mechanics (ISTM), 2022.
 - [173] D. Schanz, A. Schröder, S. Gesemann, D. Michaelis, and B. Wieneke. ‘shake the box’: A highly efficient and accurate tomographic particle tracking velocimetry (tomo-ptv) method using prediction of particle positions. In *PIV13; 10th International Symposium on Particle Image Velocimetry - ISPIV. Conference proceeding*, pages 1–13, Jan 2013.
 - [174] D. Schanz, S. Gesemann, and A. Schröder. Shake-the-box: Lagrangian particle tracking at high particle image densities. *Exp Fluids*, 57(70), 2016. DOI 10.1007/s00348-016-2157-1. URL <https://doi.org/10.1007/s00348-016-2157-1>.
 - [175] 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. DOI 10.1007/978-3-642-15825-4_10. URL https://doi.org/10.1007/978-3-642-15825-4_10.
 - [176] N. Semidetnov and C. Tropea. Conversion relationships for multidimensional particle sizing techniques. *Measurement Science and Technology*, 15:112–118, 2003. DOI 10.1088/0957-0233/15/1/015. URL <https://doi.org/10.1088/0957-0233/15/1/015>.
 - [177] D. Seol and S. Socolofsky. Vector post-processing algorithm for phase discrimination of two-phase PIV. *Experiments in Fluids*, 45:223–239, 2008. DOI 10.1007/s00348-008-0473-9. URL <https://doi.org/10.1007/s00348-008-0473-9>.
 - [178] J. D. Seymour, J. E. Maneval, K. L. McCarthy, M. J. McCarthy, and R. L. Powell. NMR velocity phase encoded measurements of fibrous suspensions. *Physics of Fluids A: Fluid Dynamics*, 5(11):3010–3012, Nov 1993. ISSN 0899-8213. DOI 10.1063/1.858709. URL <https://doi.org/10.1063/1.858709>.
 - [179] H. Shen, S. Coëtmellec, G. Gréhan, and M. Brunel. Interferometric laser imaging for droplet sizing revisited: elaboration of transfer matrix models for the description of complete systems. *Applied Optics*, 51:5357–5368, Jul 2012. DOI 10.1364/AO.51.005357. URL <https://doi.org/10.1364/AO.51.005357>.
 - [180] H. Shen, S. Coëtmellec, and M. Brunel. Simultaneous 3d location and size measurement of spherical bubbles using cylindrical interferometric out-of-focus imaging. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 131:153–159, 2013. ISSN 0022-4073. DOI 10.1016/j.jqsrt.2013.04.009. URL <https://doi.org/10.1016/j.jqsrt.2013.04.009>. Concepts in electromagnetic scattering for particulate-systems characterization.

- [181] H. Shen, S. Saengkaew, G. Gréhan, S. Coetmellec, and M. Brunel. Interferometric out-of-focus imaging for the 3d tracking of spherical bubbles in a cylindrical channel. *Optics Communications*, 320:156–161, Jan 2014. DOI <https://doi.org/10.1016/j.optcom.2014.01.020>. URL [10.1016/j.optcom.2014.01.020](https://doi.org/10.1016/j.optcom.2014.01.020).
- [182] J. Shen and H. Wang. Calculation of debye series expansion of light scattering. *Applied Optics*, 49:2422–2428, May 2010. DOI [10.1364/AO.49.002422](https://doi.org/10.1364/AO.49.002422). URL <https://doi.org/10.1364/AO.49.002422>.
- [183] J.-C. Shyu, C.-S. Wei, C.-J. Lee, and C.-C. Wang. Investigation of bubble effect in microfluidic fuel cells by a simplified microfluidic reactor. *Applied Thermal Engineering*, 30(13):1863–1871, 2010. ISSN 1359-4311. DOI [10.1016/j.applthermaleng.2010.04.029](https://doi.org/10.1016/j.applthermaleng.2010.04.029). URL <https://doi.org/10.1016/j.applthermaleng.2010.04.029>.
- [184] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition, 2015. URL <https://doi.org/10.48550/arXiv.1409.1556>.
- [185] F. Soulez, L. Denis, C. Fournier, and É. Thiébaud. Inverse-problem approach for particle digital holography: Accurate location based on local optimization. *Journal of the Optical Society of America. A*, 24:1164–1171, Apr 2007. DOI [10.1364/JOSAA.24.001164](https://doi.org/10.1364/JOSAA.24.001164). URL <https://doi.org/10.1364/JOSAA.24.001164>.
- [186] A. Srivastava, L. Valkov, C. Russell, M. U. Gutmann, and C. Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. DOI [10.48550/arXiv.1705.07761](https://doi.org/10.48550/arXiv.1705.07761). URL <https://doi.org/10.48550/arXiv.1705.07761>.
- [187] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and S. Ruslan. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- [188] J. Sun, H. Zhang, X. Wang, D. Jia, T. Liu, and M. Zhai. Determination of the orientation of transparent spheroids using interference technology. *Optics Express*, 26(11):14097–14107, May 2018. DOI [10.1364/OE.26.014097](https://doi.org/10.1364/OE.26.014097). URL <https://doi.org/10.1364/OE.26.014097>.
- [189] J. Sun, H. Zhang, W. Fan, and S. Chen. Comparison of aspect ratios of ellipsoidal particles through interferometric out-of-focus images. *Journal of the Optical Society of America A*, 38(3):395–400, Mar 2021. DOI [10.1364/JOSAA.414120](https://doi.org/10.1364/JOSAA.414120). URL <https://doi.org/10.1364/JOSAA.414120>.
- [190] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press Cambridge, Massachusetts London, England, 2014. ISBN 9780262039246.
- [191] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. DOI [10.1109/CVPR.2015.7298594](https://doi.org/10.1109/CVPR.2015.7298594). URL <https://doi.org/10.1109/CVPR.2015.7298594>.
- [192] B. Tan and D. B. Fielding. Cloud atlas: navigating the multiphase landscape of tempestuous galactic winds. *Monthly Notices of the Royal Astronomical Society*, 527(4):9683–9714, Dec 2023. ISSN 0035-8711. DOI [10.1093/mnras/stad3793](https://doi.org/10.1093/mnras/stad3793). URL <https://doi.org/10.1093/mnras/stad3793>.

-
- [193] A. B. Tayler, D. J. Holland, A. J. Sederman, and L. F. Gladden. Exploring the origins of turbulence in multiphase flow using compressed sensing mri. *Physical Review Letters*, 108: 264505, Jun 2012. DOI 10.1103/PhysRevLett.108.264505. URL <https://doi.org/10.1103/PhysRevLett.108.264505>.
 - [194] B. J. Thompson. Holographic particle sizing techniques. *Journal of Physics E: Scientific Instruments*, 7(10):781, Oct 1974. DOI 10.1088/0022-3735/7/10/001. URL <https://dx.doi.org/10.1088/0022-3735/7/10/001>.
 - [195] W. Tollmien, H. Schlichting, and H. Görtler. *Ludwig Prandtl Gesammelte Abhandlungen zur angewandten Mechanik, Hydro- und Aerodynamik*. Springer Berlin, Heidelberg, Nov 1961. DOI 10.1007/978-3-662-11836-8. URL <https://doi.org/10.1007/978-3-662-11836-8>.
 - [196] 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, 2023. URL <https://arxiv.org/abs/2302.13971>.
 - [197] C. Tropea, A. L. Yarin, and J. F. Foss. *Springer Handbook of Experimental Fluid Mechanics*. Springer Berlin, Heidelberg, 2007. DOI 10.1007/978-3-540-30299-5. URL <https://doi.org/10.1007/978-3-540-30299-5>.
 - [198] G. A. Tyler and B. J. Thompson. Fraunhofer holography applied to particle size analysis a reassessment. *Optica Acta: International Journal of Optics*, 23:685–700, 1976. DOI 10.1080/713819370. URL <https://doi.org/10.1080/713819370>.
 - [199] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders. Selective search for object recognition. *International Journal of Computer Vision*, 104:154–171, 2013. DOI 10.1007/s11263-013-0620-5. URL <https://doi.org/10.1007/s11263-013-0620-5>.
 - [200] H. C. van de Hulst and V. Twersky. *Light Scattering by Small Particles: volume 10*. Wiley Series in Pure and Applied Optics. John Wiley & Sons, New York, 1957 // 1958. ISBN 9780486642284. Reprinted by Courier Corporation, 2012.
 - [201] H. C. van de Hulst and R. T. Wang. Glare points. *Applied optics*, 30(33):4755–4763, 1991. DOI <https://doi.org/10.1364/AO.30.004755>. URL 10.1364/AO.30.004755.
 - [202] J. D. Vande Hey. *Theory of Lidar*, pages 23–41. Springer International Publishing, Cham, 2015. ISBN 978-3-319-12613-5. DOI 10.1007/978-3-319-12613-5_2. URL https://doi.org/10.1007/978-3-319-12613-5_2.
 - [203] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Neural Information Processing Systems (NeurIPS) conference*, Jun 2017. DOI 10.48550/arXiv.1706.03762. URL <https://doi.org/10.48550/arXiv.1706.03762>.
 - [204] M. Vijaymeena and K. Kavitha. A survey on similarity measures in text mining. *Machine Learning and Applications: An International Journal*, 3(2):19–28, 2016. DOI 10.5121/mlaij.2016.3103. URL <http://dx.doi.org/10.5121/mlaij.2016.3103>.
 - [205] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11:3371–3408, Dec 2010.

- [206] C.-Y. Wang, H.-Y. M. Liao, Y.-H. Wu, P.-Y. Chen, J.-W. Hsieh, and I.-H. Yeh. Cspnet: A new backbone that can enhance learning capability of cnn. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1571–1580. IEEE, 2020. DOI 10.1109/CVPRW50498.2020.00203. URL <https://doi.org/10.1109/CVPRW50498.2020.00203>.
- [207] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Jun 2020.
- [208] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. DOI 10.1109/TIP.2003.819861. URL <https://doi.org/10.1109/TIP.2003.819861>.
- [209] J. Weisman and S. Kang. Flow pattern transitions in vertical and upwardly inclined lines. *International Journal of Multiphase Flow*, 7(3):271–291, 1981. ISSN 0301-9322. DOI 10.1016/0301-9322(81)90022-7. URL [https://doi.org/10.1016/0301-9322\(81\)90022-7](https://doi.org/10.1016/0301-9322(81)90022-7).
- [210] J. Weisman, D. Duncan, J. Gibson, and T. Crawford. Effects of fluid properties and pipe diameter on two-phase flow patterns in horizontal lines. *International Journal of Multiphase Flow*, 5(6):437–462, 1979. ISSN 0301-9322. DOI 0301-9322(79)90031-4. URL [https://doi.org/10.1016/0301-9322\(79\)90031-4](https://doi.org/10.1016/0301-9322(79)90031-4).
- [211] E. W. Weisstein. Circle-circle intersection. <https://mathworld.wolfram.com/Circle-CircleIntersection.html>, 2025. From MathWorld—A Wolfram Web Resource, last accessed 15th Jun 2025.
- [212] B. Wen, Y. wu, Z. Zhuo, and X. Wu. Characterization and verification of astigmatic interferometric particle imaging for volumetric droplet 3d position and size measurement. *Advanced Powder Technology*, 32, Aug 2021. DOI 10.1016/j.appt.2021.08.006. URL <https://doi.org/10.1016/j.appt.2021.08.006>.
- [213] J. Wen and M. Breazeale. Gaussian beam functions as a base function set for acoustical field calculations. In *IEEE 1987 Ultrasonics Symposium*, pages 1137–1140, 1987. DOI 10.1109/ULTSYM.1987.199133. URL <https://doi.org/10.1109/ULTSYM.1987.199133>.
- [214] J. J. Wen and M. A. Breazeale. A diffraction beam field expressed as the superposition of gaussian beams. *Journal of the Acoustical Society of America*, 83:1752–1756, 1988. DOI DOI:10.1121/1.396508. URL <https://doi.org/DOI:10.1121/1.396508>.
- [215] C. E. Willert and M. Gharib. Three-dimensional particle imaging with a single camera. *Experiments in Fluids*, 12(6):353–358, 1992. ISSN 0723-4864. DOI 10.1007/BF00193880. URL <http://dx.doi.org/10.1007/BF00193880>.
- [216] W. J. Wiscombe. Improved mie scattering algorithms. *Applied Optics*, 19(9):1505–1509, May 1980. DOI 10.1364/AO.19.001505. URL <https://doi.org/10.1364/AO.19.001505>.
- [217] W. J. Wiscombe. Improved mie scattering algorithms. *Applied Optics*, 19(9):1505–1509, May 1980. DOI 10.1364/AO.19.001505. URL <https://doi.org/10.1364/AO.19.001505>.
- [218] R. Xu, Z. Huang, W. Gong, W. Zhou, and C. Tropea. Depth from defocus technique for irregular particle images. *Measurement*, 238:115156, 2024. ISSN 0263-2241. DOI 10.1016/j.measurement.2024.115156. URL <https://doi.org/10.1016/j.measurement.2024.115156>.

- [219] G. Yadigaroglu and G. F. Hewitt. *Introduction to Multiphase Flow: Basic Concepts, Applications and Modelling*. Springer Nature, 2018. ISBN 978-3-319-58717-2. DOI 10.1007/978-3-319-58718-9. URL <https://doi.org/10.1007/978-3-319-58718-9>.
- [220] Y. Yeh and H. Z. Cummins. Localized fluid flow measurements with an he-ne laser spectrometer. *Applied Physics Letters*, 4(10):176–178, May 1964. ISSN 0003-6951. DOI 10.1063/1.1753925. URL <https://doi.org/10.1063/1.1753925>.
- [221] H. T. Yura and S. G. Hanson. Optical beam wave propagation through complex optical systems. *Journal of The Optical Society of America A-optics Image Science and Vision*, 4(10):1931–1948, 1987. DOI 10.1364/JOSAA.4.001931. URL <https://doi.org/10.1364/JOSAA.4.001931>.
- [222] H. Zhang, X. Wang, J. Sun, D. Jia, and T. Liu. Multidispersed bubble-size measurements by interferometric particle imaging at scattering angles of 90° and 45°. *Applied Optics*, 57(36):10496–10504, Dec 2018. DOI 10.1364/AO.57.010496. URL <https://doi.org/10.1364/AO.57.010496>.
- [223] X. Zhang, Z. Dong, H. Wang, X. Sha, W. Wang, X. Su, Z. Hu, and S. Yang. 3d positioning and autofocus of the particle field based on the depth-from-defocus method and the deep networks. *Machine Learning: Science and Technology*, 4(2):025030, Jun 2023. DOI 10.1088/2632-2153/acdb2e. URL <https://dx.doi.org/10.1088/2632-2153/acdb2e>.

List of Conference Contributions

- [LSBK22] Robin Leister, Christian Sax, Andreas Brümmer, and Jochen Kriegseis. Laser-optical shear-flow analysis across the annular gap of a simplified displacement compressor model. Presentation given on the 11th International Conference on Screw Machines (2022), Dortmund, Germany, 7th-8th of September 2022, 2022.
- [LSBK24] Heinrich Lange, Christian Sax, Andreas Brümmer, and Jochen Kriegseis. Towards dptv and ipi across curved displacement-compressor surfaces: Impact of astigmatism on particle-image characteristics of bubbles and tracers. Presentation given on the 1st European Fluid Dynamics Conference (EFDC1 2024), Aachen, Germany, 16th-20th of September 2024, 2024.
- [SDK23] Christian Sax, Maximilian Dreisbach, and Jochen Kriegseis. Blasengrößenbestimmung durch interferometric particle imaging (ipi) in vorwärts-, seitwärts- und rückstreuung. Presentation given on the 30. Fachtagung "Experimentelle Strömungsmechanik" (2023), München, Germany, 5th-7th September 2023, 2023.
- [SGK25] Christian Sax, Roland Griesmaier, and Jochen Kriegseis. A novel diagnostics tool for the experimental characterization of dispersed multi-phase flows. Presentation given on the 12th International Conference on Multiphase Flow (ICMF 2023), Toulouse, France, 12th-16th of May 2025, 2025.
- [SK23] Christian Sax and Jochen Kriegseis. Simulation of 3d interferometric particle imaging experiments for the measurement of multi-phase flows. Presentation given on the 2nd MathSEE Symposium, Karlsruhe, Germany, 27th-29th of September 2023, 2023.

- [SLBK23] Christian Sax, Robin Leister, Andreas Brümmer, and Jochen Kriegseis. Towards interferometric particle imaging for bubble-size estimation in thin annular gap-flows. Presentation given on the 11th International Conference on Multiphase Flow (ICMF 2023), Kobe, Japan, 2nd-7th of April 2023, 2023.

List of Journal Publications

- [LSBK24] Heinrich Lange, Christian Sax, Andreas Brümmer, and Jochen Kriegseis. Astigmatism quantification for depth localization of bubbles and tracers across curved surfaces. *IOP Conference Series: Materials Science and Engineering*, 1322, 2024.
- [SDK25] Christian Sax, Maximilian Dreisbach, and Jochen Kriegseis. Interferometric particle imaging for particle sizing in the front-, side-, and backscatter region - towards single optical access ipi. *Physical Review Applied*, Manuscript under Review, Jun 2025. Preprint available <https://doi.org/10.48550/arXiv.2303.16013>. Cited version from the 26th of June 2025.
- [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, Sep 2023.
- [SGK25] Christian Sax, Roland Griesmaier, and Jochen Kriegseis. Particle localization with dptv and sizing with ipi using a forward model and optimization. *Measurement Science and Technology*, Manuscript under Review. Cited version from the 26th of June 2025, 2025.
- [SK25a] Christian Sax and Jochen Kriegseis. On the particle image overlap in single camera defocusing approaches, Jun 2025. Published Preprint <https://doi.org/10.48550/arXiv.2506.18170>.
- [SK25b] Christian Sax and Jochen Kriegseis. Pattern-based phase-separation of tracer and dispersed phase particles in two-phase defocusing particle tracking velocimetry, Jun 2025. Published Preprint <https://doi.org/10.48550/arXiv.2506.18157>.
- [SSL⁺23] C. Sax, A. Stroh, R. Leister, M. Dreisbach, J. Kriegseis, C. Denda, and P. Bürk. Fluid-mechanical evaluation of different clutch geometries based on experimental and numerical investigations. *Forschung im Ingenieurwesen*, 87:1297–1306, 2023.

Acronyms and symbols

Acronyms

1D	One-Dimensional
2D	Two-Dimensional
2D-2C	Two-Dimensional Two-Components
3C	Three components
3D	Three-Dimensional
3D-3C	Three-Dimensional Three-Components
Adam	Adaptive Moment estimation
ADB	Alexander's Dark Band
AP	Average Precision
APT_V	Astigmatism Particle Tracking Velocimetry
BB	Bounding Box
CA	Class Accuracy
CCD	Charge-Coupled Device
CNN	Convolutional Neural Network
CSP	Cross Stage Partial
DEHS	Di-Ethyl-Hexyl Sebacat
DPT_V	Defocusing Particle Tracking Velocimetry
DSE	Debye Series Expansion
DTPF	Dispersed Two-Phase Flow
EIT	Electrical Impedance Tomography
FC	Fully Connected
FFT	Fast Fourier Transform
FM	Forward Model
FN	False Negative
FOV	Field of View
FP	False Positive

FPN	Feature Pyramid Network
FRA	Free Remaining Area
GAN	Generative Adversarial Network
GO	Geometrical Optics
GP	Glare Point
GPI	Glare Point Image
GT	Ground Truth
ILIDS	Interferometric Laser Imaging for Droplet Sizing
IP	Inverse Problem
IPI	Interferometric Particle Imaging
IoA	Intersection over Area
IoM	Intersection over Minimum
IoU	Intersection Over Union
LDV	Laser Doppler Velocimetry
LIDAR	Light Detection and Ranging
LIF	Laser-Induced Fluorescence
LMT	Lorenz-Mie Theory
MLP	Multi-Layer Perceptron
MRI	Magnetic Resonance Imaging
MS-COCO	Microsoft Common Objects in Context (dataset)
MSE	Mean Squared Error
OCT	Optical Coherence Tomography
OLR	Overlap Ratio
PANet	Path Aggregation Network
PDF	Probability Density Function
PI	Particle Image
PIV	Particle Image Velocimetry
PTV	Particle Tracking Velocimetry
px	Pixels
R-CNN	Region Based Convolutional Neural Network
ReLU	Rectified Linear unit
RGB	Red Green Blue
RoI	Region of Interest

RPN	Region Proposal Network
RQ	Research Question
RTM	Ray Transfer Matrix
SGD	Stochastic Gradient Descent
SGDM	Stochastic Gradient Descent with momentum
SNR	Signal-to-Noise Ratio
SPP	Spatial Pyramid Pooling
SSIM	Structural Similarity Index
Stk	Stokes number
TAP	Truncated Average Precision
TCA	Total Covered Area
TE	Transverse Electric (perpendicular polarisation)
TM	Transverse Magnetic (parallel polarisation)
TN	True Negative
TP	True Positive
UPDV	Ultrasonic Pulsed Doppler Velocimetry
YOLO	You Only Look Once

Latin letter -Upper case

$A_1^{x,y}$	first upper entry of the RTM $M_1^{x,y}$ from the particle to the principal plane
$A_2^{x,y}$	first upper entry of the RTM $M_2^{x,y}$ from the principal plane to the image plane
A_{GP}	complex amplitude of of a wave along the particles projection
A_{PI}	particle image area
A_{image}	image area
A_{tot}	first upper entry in the RTM
$B_1^{x,y}$	second upper entry of the RTM $M_1^{x,y}$ from the particle to the principal plane
$B_2^{x,y}$	second entry of the RTM $M_2^{x,y}$ from the principal plane to the image plane
B_{tot}	second upper entry in the RTM
$C_1^{x,y}$	first lower entry of the RTM $M_1^{x,y}$ from the particle to the principal plane
$C_2^{x,y}$	first lower entry of the RTM $M_2^{x,y}$ from the principal plane to the image plane
C_ℓ	ℓ^{th} class
C_{tot}	first lower entry in the RTM
$D_1^{x,y}$	second lower entry of the RTM $M_1^{x,y}$ from the particle to the principal plane

$D_2^{x,y}$	second lower entry of the RTM $M_2^{x,y}$ from the principal plane to the image plane
$D_{a,b}$	used as abbreviation term for the Riccati-Bessel function
D_a	aperture diameter
D_{pp}	peak to peak distance
D_{tot}	second lower entry in the RTM
E_ℓ	complex coefficient
F	fringe frequency
G	complex field of a wave in a plane
G_0	complex field of a wave in the source plane
G_1	complex field of a wave in the principal plane
G_2	complex field of a wave in the image plane
$G_{2,GP}$	complex field of a glare point in the image plane
G_P	complex field of a particle image
$G_{P,pos}$	complex field of a particle image with suppressed interference
G_{obs}	observed field in the image plane, i.e. the recorded image
H	Hyperparameters
I	absolute intensity
$I^{(p)}$	absolute intensity of scattering order p
I_{BG}	background intensity
I_{Mie}	absolute intensity computed by the LMT
I_{PI}	intensity of the particle image
L_{BB}	bounding box loss
L_C	class loss
L_{CE}	cross entropy loss
L_D	discriminator loss
L_G	generator loss
L_{MSE}	mean squared error loss
$L_{multitask}$	multi-task loss
$L1$	L1 loss
M_1	ray transfer matrix from particle to the principal plane
M_2	ray transfer matrix from the principal plane to the image plane
M_G	magnification
$M_{I,i \rightarrow j}$	ray transfer matrix for an interface with the direction of travel from medium i to j

M_{lens}	ray transfer matrix for a lens
$M_{\text{p},i}$	ray transfer matrix for propagation in medium i
M_{tot}	ray transfer matrix
N	number
N_{C}	number of classes
N_{noOL}	number of particle images without overlap
N_{OL}	number of particle image overlaps
N_{P}	number of particles
N_{S}	number of stripes
N_{px}	number of pixels
Q_{ℓ}	complex coefficient
R^{11}	reflection coefficient from medium one back to medium one
R^{22}	reflection coefficient from medium two back to medium two
$R_{1,2}$	reflection coefficient with the index one denoting the TE and index two the TM mode
R_{I}	radius of a curved interface
S_1	complex amplitude of the scattered wave in TE mode
S_2	complex amplitude of the scattered wave in TM mode
S_{GP}	complex amplitude of a glare point
$T(\xi, \eta)$	transmission function of the aperture
T^{12}	transmission coefficient from medium one into medium two
T^{21}	transmission coefficient from medium two into medium one
$T_{1,2}$	transmission coefficient with the index one denoting the TE and index two the TM mode
U	scalar field of a wave
V	visibility
V_{C1}	first condition of the visibility
V_{C2}	second condition of the visibility
$\{X, Y, Z\}$	coordinate system
\mathcal{D}	discriminator
\mathcal{G}	generator
$\mathcal{M}_{\text{circle}}$	circle mask
$\mathcal{M}_{\text{comb}}$	combined circle mask

\mathcal{P}	mismatch functional
\mathcal{R}	recall
\mathcal{S}	seeding density
Pr	precision
Se	sensitivity of the glare point spacing

Latin letter - Lower case

(u, v)	pixel coordinates
(x_2, y_2, z_2)	point in the image plane
(ξ, η)	coordinates in the principal plane
$a(z)$	activation function of a neuron
a_ℓ, b_ℓ	Mie coefficients
b	bias term of a neuron
c	constant
d_C	distance between circle centres
d_{PI}	particle image diameter
d_P	particle diameter
d_s	distance from the slit to the screen in z -direction in Young's fringe experiment
e	Euler number
$f(x)$	generic function
$f_\#$	aperture f-number (focal length divided by the aperture diameter)
f_{lens}	focal length of a lens
h	hidden layer block, hidden filter kernel for a convolutional layer
i	imaginary number ($i = \sqrt{-1}$)
k	used as index
l	effective optical length a wave travels at a particle
l_c	characteristic length scale
m	relative refractive index
n	refractive index
n_1	refractive index of the inside of a particle
n_2	refractive index of the surrounding medium
n_a	refractive index of air
n_g	refractive index of glass

n_w	refractive index of water
p	scattering order
p_C	probability of a class
q	constant
r_{PI}	particle image radius
s_0	distance imaging optics to focal plane
t	time
t_F	response time fluid
t_P	response time particle
u_c	characteristic velocity of fluid
w	relative coordinate of glare point on the sphere projection
$w^{(p)}$	relative coordinate of glare point of order p on the sphere projection
w^{phys}	absolute coordinate of glare point in meter
w_p^{phys}	absolute position of a glare point of order p in meter
w_x	relative x -coordinate of glare point on the sphere projection
w_y	relative y -coordinate of glare point on the sphere projection
x_{GP}	x -position of a glare point in the image
x_P	x -position of a particle
x_{PI}	x -position of a particle image
x_{noise}	latent noise vector
x_{ray}	distance of a ray to the optical axis
x_{real}	set of real data
x_{size}	dimensionless size parameter of sphere
y_{GP}	y -position of a glare point in the image
y_P	y -position of a particle
y_{PI}	y -position of a particle image
z	generic argument
z^{eff}	effective optical distance (distance divided by the refractive index)
z_{12P}^{defoc}	distance from the focused image plane to the principal plane for a particle outside of the object side focal plane
z_{12P}^{foc}	distance from the focused image plane to the principal plane for a particle in the object side focal plane
z_P	particles distance to the focal plane in z -direction

z_{P2F}	distance from the principal plane to the focal plane
$z_{\text{P,lens}}$	distance from the particle to the lens / principal plane along the z axis
z_{a}	length in air
z_{aperture}	z position of the aperture from the particle
z_{g}	length in glass
z_{im}	distance from the lens to the camera chip
z_{obj}	object side distance from the lens to the particle
z_{ray}	z -position of a ray along the optical axis
z_{w}	length in water
ℓ	used as an index
ℓ_{stop}	stopping index for the order of partial waves in the LMT

Greek letters - Upper case

Δ_{GP}	relative glare point distance in w coordinates
$\Delta_{\text{GP}}^{\text{Err}}$	erroneous relative glare point distance in w coordinates
$\Delta_{\text{GP}}^{\text{phys}}$	absolute glare point distance in meter
Ω_{PI}	set of pixels belonging to a particle image

Greek letters - Lower case

(ξ, η)	coordinate system in the principal plane
α_{AP}	opening angle of the aperture
β_{i}	angle of incidence
$\beta_{\text{i,B}}$	angle of incidence at the Brewster condition
β_{r}	reflected angle
β_{t}	angle of transmission
κ	wave number
λ	wavelength of light
λ_0	wavelength of light in vacuum
μ	mean
μ_{P}	absolute viscosity of fluid
ϕ_{ray}	angle of a ray to the optical axis
ψ_{I}	parameter vector for multiple particles in an image
ψ_{P}	parameter vector of a single particle

ρ_P	physical density of a particle
σ	standard deviation
θ	scattering angle
θ_0	scattering angle at the middle of the particle ($w = 0$)
θ_{tot}	total deflection angle
φ	phase of a wave
ϱ	3D vector
ω	weight of a neuron

Mathematical operators

$(\dots)'$	derivative w.r.t. the argument
$(\dots)^T$	transposed vector
$(\dots)_\ell$	index used in sums or products
$(\dots)^{-}$	average
\cap	intersection
\cup	union
Δ	difference
$(\hat{\dots})$	ground truth
$\Im\{\dots\}$	imaginary part
\mathcal{F}	Fourier transform
\mathcal{F}_2	2D Fourier transform
∇	Laplace operator
$\partial(\dots)$	partial derivative
$\Re\{\dots\}$	real part
$d(\dots)$	differential element
\star	normalised cross-correlation

Special functions

$\mathbf{1}(z)$	indicator function
$\delta(z)$	Dirac function
$\Gamma_\ell(z), \Upsilon_\ell(z)$	Riccati-Bessel functions of third kind and order ℓ
\mathcal{H}_ℓ	Hankel function of order ℓ
$\mathcal{J}_\ell(z)$	spherical Bessel function of first kind and order ℓ

$\mathcal{Y}_\ell(z)$	spherical Bessel function of second kind and order ℓ
$P_\ell^{(j)}$	associated Legendre polynomial of order j and degree ℓ
Π_ℓ, τ_ℓ	Mie angular functions

List of Figures

2.1	Flow patterns in a vertical pipe. The flow patterns are also marked by whether they classify as a disperse or partially disperse partially continuous two-phase flow. Figure based on [21].	9
2.2	Flow patterns in a horizontal pipe. The flow patterns are marked by whether they classify as a disperse, partially disperse partially continuous or continuous two-phase flow. Figure based on [21].	10
2.3	Defocusing approach visualised. Shown are two particles, one closer to the focal plane (blue) resulting in a small defocus circle and a particle further away from the focal plane (red) resulting in a larger defocus circle. The image plane (i.e. the camera chip) and the (object side) focal plane are shown. Additionally, is the image side focal plane shown, which is the plane at which the defocused particle would be imaged in focus. A 3D close-up on how the defocused particle image changes between the image side focal plane and the actual image plane is shown. Modified figure based on [113].	14
2.4	"Working Principle of DPTV visualised. As the particle moves away from the focal plane the particle image visible on the image plane becomes defocused and the diameter of the PI increases." [SGK25]. Figure from [SGK25]	15
2.5	Resolution of a flow field, based on the number of particles in the fluid visualised. A flow field of a vortex is shown as an example. The left image contains $N_P = 20$ particles and the right image contains 60 particles. The effect of defocusing and the seeding density on the particle image overlap is shown.	16
2.6	Light scattering on a spherical particle as visualised by the geometrical optics model. The relationship between glare points and the scattering angle θ is illustrated. The scattering of the dominant light rays is indicated on the blue spheres. The positions of the glare points, as they would appear to an observer at the angle θ , can be seen on the black spheres. Figure adapted from [SDK25].	19
2.7	"A bubble $m < 1$ (left) and a droplet $m > 1$ (right) are shown. Depicted are the reflected ray $p = 0$ (●) and the first three orders of refracted rays $p = 1$ (●), $p = 2$ (●), and $p = 3$ (●). The position of the glare point on the sphere's surface is indicated by the w coordinate. Note that the light is monochromatic and the colours encode the scattering order." [SDK25]. Figure adapted from [SDK25]	19
2.8	"(a) Light paths of the $p = 0$ and $p = 1$ rays through the particle using the GO model to determine the position of the glare point, on the example of a droplet. The working principle of IPI is visualized in (b). From the observation angle θ , two glare points are visible on the particle surface, when imaged in focus. The images of the glare points begin to overlap and form an interference pattern as they move away from the focal plane. They eventually collapse into a single circle at a sufficient distance from the focal plane. In addition, the analogy of the glare point spacing (particle size) in IPI to a double slit in Young's fringe experiment [73] is visualized ((b) top). The intensity maxima are depicted as dark stripes for better visibility (inverted intensity)." [SDK25]. Figure adapted from [SDK25].	20

2.9	A light ray (scattering order $p = 1$) interacting with a sphere of refractive index n_1 in a surrounding of refractive index n_2 (left). At every interaction of the ray with the interface is a part of the ray transmitted and a part reflected. The transmission and reflection coefficients at each interaction with an interface are shown, with e.g. T^{12} denoting a transmission from medium one (n_1) to medium two (n_2) and R^{11} denoting a reflection from medium one back to medium one. The different light intensity that is observed based on the scattering angle is visualised in the second figure (right).	23
2.10	Mie angular functions $\Pi_\ell(\theta)$ (top) and $\tau_\ell(\theta)$ (bottom) visualised for the partial waves of order $\ell = 2, \dots, 6$	25
2.11	Logarithmic intensity of light scattered by an air bubble in water over the scattering angle θ , according to the Mie solution. The depicted bubble has a diameter $d_P = 20 \mu\text{m}$, $n_1 = 1$, $n_2 = 1.333$ and the light has a wavelength of $\lambda = 532 \text{ nm}$, ($m = 1/1.333$, $x_{\text{size}} = 314.9$).	25
2.12	Logarithmic intensity of light scattered by an air bubble in water over the scattering angle θ . The intensity of the first three scattering orders ($p = 0, 1, 2$) in TM-mode, according to the DSE of the LMT are shown. All intensities are scaled identically to visualise the intensity loss of higher scattering orders. The depicted bubble has a diameter $d_P = 20 \mu\text{m}$, $n_1 = 1$, $n_2 = 1.333$ and the light has a wavelength of $\lambda = 532 \text{ nm}$, ($m = 1/1.333$, $x_{\text{size}} = 314.9$).	27
2.13	Notation for the ray transfer matrix system visualised. A ray is characterised by its distance and it's angle to the optical axis	29
2.14	Schematic structure of a fully connected feed forward neural network. There is an input layer x_i , two hidden layers h_i and an output layer y_i . Figure adapted from [172]	33
2.15	Schematic structure of a neuron in a neural network. The neuron receives inputs x_i which are weighted by the trainable weights ω_i . Each input is weighted and then summed up and a bias b is added. The results is then fed through an activation function $a(z)$ to produce the final output of the neuron y . Figure adapted from [172].	34
2.16	Spatial convolution visualised: A 3×3 kernel (gray cells) is moved over a 5×5 input (green cells) with padding (white cells). The kernel moves with a stride of one producing a 5×5 output (blue cells). Figure adapted from [49].	39
2.17	Simple architecture of a CNN classifier. An input image is processed using a convolutional layer with multiple filters. The network used multiple convolutional layers, decreasing in size but increasing in depth, increasing the level of abstraction of features. The final feature vector is flattened and processed by a two fully connected layers to perform the classification tasks. Figure adapted from [130]	40
2.18	Residual block: The output from a previous layer is added to the output of the current layer. Figure adapted from [72]	41
2.19	Structure of Faster R-CNN: The backbone, the RPN and the detection head are marked, as well as their tasks are pointed out. Modified figure from [30].	43
3.1	The relationship between glare point distance and the resulting interference patterns is visualised. As the glare point distance increases, the fringe frequency correspondingly increases. While the angular frequency remains constant along the z -axis, the interference pattern progressively spreads out with increasing distance from the focal plane.	48
3.2	"The concept of visibility and the two conditions are visualized in an example at a single scattering angle. Shown is the intensity over the projected particle surface as determined by the LMT (middle) and the decomposition into individual scattering orders with the Debye-Series expansion (left, right). The glare points are the peaks of intensity." [SDK25]. Figure from [SDK25].	50

- 3.3 "The complex amplitude S_2^2 of each scattering order is plotted over all scattering angles. Depicted are the Debye plots for (a) a water droplet in air ($m = 1.333 + i1.82E - 09$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1476.3$) and (b) an air bubble in water ($m = 1/1.333 + i$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1967.9$). The visibility criteria are visualized: Scattering angles satisfying both criteria are marked green (●), whereas angles only matching the first condition but not the second one are marked red (●)."[SDK25]. Figure from [SDK25] 51
- 3.4 "Visibility over the scattering angle for a water droplet in air ($m = 1.333 + i1.82 \cdot 10^{-9}$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1476.3$) in (a) TE- (V_1) and (b) TM-mode (V_2), and for an air bubble in water ($m = 1/1.333 + i$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1967.9$) in (c) TE- (V_1) and (d) TM-mode (V_2). A sufficient scattering angle for IPI is given whenever the visibility plot has a value close to one. The glare points parings (0,1), (0,2), (0,3) and (1,2) are shown. In addition, the rainbow angles of the $p = 2$ and $p = 3$ scattering orders and Alexander's dark band (ADB) are marked in (a) and (b). The Brewster angle is marked in (d)."[SDK25]. Figure from [SDK25] 52
- 3.5 "Intensity over the particle surface (the glare points are the peaks): The scattering intensity $A_{\text{GP}}(w)$ is plotted over the particle surface w for a water droplet in air ($d_p = 250 \mu\text{m}$, $\lambda = 532 \text{ nm}$, $\Rightarrow x_{\text{size}} = 1476.3$ and $m = 1.333 + i1.82 \cdot 10^{-9}$). The scattering uncertainty $\Delta\theta = 3.7^\circ$ was chosen based on the optical system used for the validation experiment. The scattering angle θ_0 around which $A_{\text{GP}}(w)$ was calculated is shown at the top of each plot. The intensities of the scattering orders $p = 0$ (●), $p = 1$ (●), $p = 2$ (●) and $p = 3$ (●), as well as the integral intensity of all scattering orders (Mie-intensity) (●) are shown. Unless indicated all plots are on the same intensity scale."[SDK25]. Figure from [SDK25] 54
- 3.6 "Intensity over the particle surface (the glare points are the peaks): The scattering intensity ($A_{\text{GP}}(w)$) is plotted over the particle surface w for an air bubble in water ($d_p = 250 \mu\text{m}$, $\lambda = 532 \text{ nm}$, $\Rightarrow x_{\text{size}} = 1967.9$ and $m = 1.333 + i1.82 \cdot 10^{-9}$). The scattering uncertainty $\Delta\theta = 3.7$ was chosen based on the optical system used for the validation experiment. The scattering angle θ_0 around which $A_{\text{GP}}(w)$ was calculated is shown at the top of each plot. The intensities of the scattering orders $p = 0$ (●), $p = 1$ (●), $p = 2$ (●) and $p = 3$ (●), as well as the integral intensity of all scattering orders (Mie-intensity) (●) are shown. Unless indicated all plots are on the same intensity scale."[SDK25]. Figure from [SDK25] 55
- 3.7 "Glare point maps: Glare point position w over the scattering angle visualized for the reflected and first three orders of refracted light rays for a droplet (a) and and a bubble (b). For the droplet in (a), the rainbow angles and ADB are marked. For the bubble in (b) it is shown how to measure the glare point spacing from the plots."[SDK25]. Figure from [SDK25] 56
- 3.8 "Interrelationship between intensity, scattering angle θ and glare point position w . Glare point plots represent vertical slices placed on the glare point maps (horizontal slice)."[SDK25]. Figure from [SDK25] 57
- 3.9 "The optical system of the validation experiment is shown. Depicted are a schematic principle of the experimental setup (a) and an image of the real experimental setup (b). A schematic close-up of the bubble generator (cathode) and the field of view (FOV) relative to the tip of the cathode-hook are shown (c). The physical model for the optical path from the bubble to the camera is shown up in a close-up (c)"[SDK25]. Modified figure from [SDK25] 58

- 3.10 "Experimental results of the validation experiment. Raw images from the validation experiment are shown in (a). For the image at 45° , a secondary frequency can be observed, which is the result of a third glare point ($p = 2$) in the TE-mode. It can be seen that the particles visually become less distinct from the background and the intensity decreases in the backscatter region. The determined distribution functions for all investigated scattering angles $45^\circ, 99^\circ, 116^\circ, 122^\circ, 135^\circ$ and 154° (angle in water at the bubble) are plotted in (b). The number of particles in each PDF is shown above the graph. The scattering angle of 45° the PDF was evaluated with Eqs. (2.11) ($45^\circ S$) and (3.4) (45°). The smallest measurable bubble diameter for $\theta = 45^\circ$ was reached and is marked with $d_{\min}^{45^\circ}$. The PI brightness in the raw images (a) is lower at higher scattering angles, as the overall intensity of the interference pattern is reduced. This can be observed by comparing the intensity in Fig.,3.12."[SDK25]. Modified figure from [SDK25] 60
- 3.11 "Figure (a) shows the different kernels used to average the signal along the maxima of the interference patterns and the resulting signal across the PI. The intensity in (a)(top) is inverted for better visibility. The peak-to-peak distance between the maxima of the interference pattern over the radius ($\theta = 45^\circ$) is shown in (b). Figure (c) shows the relative error (for an uncertainty of 1 px) in the measurement of the peak-to-peak distance over the particle diameter for different scattering angles (for the setup used in the validation experiment). In addition, the upper and lower limits of the measurable particle diameters at each scattering angle can be determined from the respective ends of the error lines."[SDK25]. Figure from [SDK25] 62
- 3.12 "Peak intensity of the interference pattern for a water droplet in air ($m = 1.333 + i1.82 \cdot 10^{-9}$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1476.3$) in (a) TE- and (b) TM-mode, and an air bubble in water ($m = 1/1.333 + i0$, $\lambda = 532 \mu\text{m}$, $d_p = 250 \mu\text{m}$, $\Rightarrow x_{\text{size}} = 1967.9$) in (c) TE- and (d) TM-mode."[SDK25]. Figure from [SDK25] 64
- 3.13 "The signal strength (mean intensity of the maxima in every PI) over the background intensity for the images used in the validation experiment. The values correspond to the images shown in Fig.3.10a. It can be seen that the particles gradually become less distinct from the background and the intensity decreases from the front- over the side- into the backscatter region."[SDK25]. Figure from [SDK25] 65
- 3.14 "Glare point spacing $\Delta_{\text{GP}(\theta)}$ over the scattering angle for droplets ($m = 1.333$) (a) and bubbles ($m = 1/1.333$) (b). The lines are colour coded to indicate the glare point spacing. Only the glare point spacing of the glare point paring dominant at the scattering angle is shown."[SDK25]. Figure from [SDK25] 66
- 4.1 "Scattering model (left) and its integration into the propagation model (right) visualised. Incident light is scattered at the particle. The scattering orders $p = 0$ and $p = 2$ and the position of the GPs w_0 and w_2 are shown. The distance travelled from the incident plane to the object plane is in the surrounding medium l_{out} or inside the particle l_{coord} . The GPs are then modelled as point emitters in the source plane. The length l_{out} allows both GPs to be modelled in the same plane."[SGK25]. Figure from [SGK25]. 71

4.2	"Propagation model and the incorporation of the scattering model visualised. Two particles (double emitters) are shown in a simple optical system consisting of a single lens with an aperture. The GPs of each particle are located in their respective source plane, which has a distance of z_P from the focal plane of the optical system. The field in the source plane is G_0 . The lens and aperture are located in the principal plane. The field G_1 in the principal plane is clipped by the aperture function T . The cropped field G_1T then propagates to the image plane to form the field G_2 which is the final image seen by the camera $G_{\text{image}} = G_2 $. Also shown are the coordinates of a particle in the source plane (x_P, y_P) and its counterpart (x_{PI}, y_{PI}) of the particle image in the image plane."[SGK25]. Modified figure from [SGK25].	73
4.3	"The superposition principle for the construction of a PI from the individual GPIs is visualised (top). The construction of a full camera image with multiple PIs is shown in the lower half (bottom)."[SGK25]. Figure from [SGK25]	75
4.4	"Gray value intensity distribution of the imaged PI field for different distances z_P from the focal plane as computed with the forward model. For comparison, the theoretical PI diameter according to the Olsen-Adrian model [146] (\bullet) and the GO model from Eq. (2.4) (\circ) are converted to pixels and added to the diagram."[SGK25]. Figure from [SGK25]	76
4.5	"Optical system of the experiment used to obtain the experimental PIs of the bubbles. The bubbles are located at a distance z_P from the focal plane. The focal plane is located in a water tank behind a glass wall. The imaging system is a single lens with an aperture in the principal plane. The camera chip is located in the image plane. The refractive indices of each part of the optical system are given."[SGK25]. Figure from [SGK25]	76
4.6	"Experimental (left) and simulated (right) particle images for three different bubble sizes (top). In the bottom row the signal from averaging the PI along the stripes of the pattern is shown. Below each graph the used z -position and size of each bubble is displayed. The uncertainty of the size estimation with the FFT approach is also displayed. The size of the particle was only varied within a tenth of the uncertainty of the FFT approach to obtain the simulated image."[SGK25]. Figure from [SGK25]	77
4.7	"Particle images for different particle diameters and distances from the focal plane. Larger particles require greater distances from the focal plane to be sufficiently defocused (i.e., the PI resembles a single circle). PIs without interference, as created with Eq. (4.20), are shown, to visualise the double circle nature of double emitter PIs (cp. also the blue circles). The images without interference are used to calculate $\mathcal{P}_{P,\text{pos}}$. The influence of the particle diameter in combination with the z -position on the PI shape becomes obvious, when omitting the interference. Also the effect of aliasing as a function of diameter and z -position can be observed as marked by the red circle. Recall that aliasing can occur in two ways, where either the particle becomes larger with accordingly increasing number of fringes or the PI becomes smaller, such that the remaining PI-intersection for the fringes diminishes."[SGK25]. Figure from [SGK25]	79
4.8	"Scheme of the IP approach in the context of a measurement. First images are recorded and a detection algorithm (e.g. a CNN) is used to generate BBs. The BBs are used to initialise ψ_P and the IP. In the IP ψ_P is fed into the FM to produce a simulated image. The simulated image is then compared to the observed image by the mismatch functional, which quantifies the mismatch by a single value. This value is then minimized (e.g. by a trust-region interior-point algorithm) until a convergence criterion is reached. The particle field is then reconstructed from the final iteration of ψ_P . The IP is analogue for a single or for multiple particles."[SGK25]. Modified figure from [SGK25]	81

4.9	"IP approach and FFT for the diameter measurement with known particle position (x_P, y_P, z_P). Both approaches were given the ground truth particle position and only the particle diameter was determined. For the simulated observed image, in (a) the particle z -position was varied with a fixed diameter, and in (b) the particle diameter at fixed z -position was varied. The red dashed line (---) represents the micro-meter accuracy for the present optical system. The red dot (•) shows where the two test series intersect in parameter space."[SGK25]. Figure from [SGK25]	82
4.10	"Optimization of the particle z -position with the ground truth of the other particle parameters (x_P, y_P, d_P) for a varying z -position and fixed diameter (a) and for a varying diameter with fixed z -position (b). The black dashed line (---) represents the pixel accuracy. The red dot (•) shows where the two test series intersect in parameter space."[SGK25]. Figure from [SGK25]	83
4.11	"Probability density function (PDF) of the errors of the IP approach (1 px initialisation error) for 200 single particle images with randomly distributed positions and sizes. The particles have diameters in the range of $d_P = [20, 300] \mu\text{m}$ and a random distance from the focal plane in the range of $z_P = [1, 8] \text{mm}$. The in-plane position was randomly generated for the PI to be still fully within the image. The pixel accuracy is shown as the black dashed line, and the micro-meter accuracy (for the present optical system) as a red dashed line."[SGK25]. Modified figure from [SGK25]	84
4.12	"Real PIs from the bubble experiment [SDK25] (left), the PI from the IP result (middle) and the image overlay (right) (observed - blue channel; IP - red channel). Given are the ground truth values \hat{z}_P, \hat{d}_P (i.e., measurement with conventional approach) with the measurement uncertainty (left), the values obtained from the IP approach z_P, d_P (middle) and the discrepancy of the two methods (right). The same images as in Fig. 4.6 for the validation of the FM were used. Similar to Fig. 4.6, the blue line on the right represents the experiment and the orange line, the reconstruction by the IP."[SGK25]. Figure from [SGK25]	86
4.13	"Consistency of the IP approach for images with increased noise levels. Images from the bubble experiment were used and Gaussian noise was added to decrease the SNR (top). The change of the optimization result of the IP approach for different noise levels is shown for the x -position (left), the z -position (middle) and the diameter (right). The values ($\hat{x}_P, \hat{z}_P, \hat{d}_P$) refer to the results of the IP approach without added noise, see Fig. 4.12. The error is the deviation from the 'noise free' results. The same images as in Fig. 4.6 for the validation of the FM were used."[SGK25] Figure from [SGK25]	87
4.14	"Probability density functions for the in plane x and y errors (a), the out-of-plane errors (b) and the diameter estimation error (c). For the diameter error also the FFT approach is plotted as comparison. Similarly to the FFT, the IP approach was initialised with 1 px uncertainty. Ten images with 20 particles each were generated with random particle parameters in $d_P \in [20, 300] \mu\text{m}$ and $z_P \in [1, 8] \text{mm}$. The pixel accuracy is shown as the black dashed line, and the micro-meter accuracy as a red dashed line. The yellow dashed line in (c) marks the threshold beyond which outliers are defined."[SGK25]. Figure from [SGK25]	89
4.15	"The z -position (a) and diameter (b) errors over the cumulated Intersection over area (IoA) of the PIs. The IoA is the intersection area of a PI divided by its area. The accumulated IoA is the summed IoA of a PI with all other PI overlaps and shows the amount of the PI area that is overlapped. In (a) the dashed black line represents the pixel accuracy. In (b) the thresholds of a $1\mu\text{m}$ diameter error (red line) and for the outlier definition of an error of $10\mu\text{m}$ (yellow line) are visualised. The errors are shown for the IP approach and for the FFT approach both with a 1 px BB accuracy for initialisation. The FFT approach and the IP approach are initialised from a ± 1 px BB placement accuracy."[SGK25]. Figure from [SGK25]	90

4.16	"The estimation results of two of the 10 synthetic test images visualized for the FFT + BB (1 px placement uncertainty) and the IP approach. A colored dot is added on each PI to show the error of the diameter determination: For an error $< 1 \mu\text{m}$ a green dot is placed (\bullet), for an error in $[1, 10] \mu\text{m}$ a blue dot (\bullet) and for outlier particles (error $> 10 \mu\text{m}$) a red dot is placed (\bullet)."[SGK25]. Figure from [SGK25]	91
4.17	"Synthetic observed image and the reconstruction by the IP approach visualized. The reconstructed particles are compared to the ground truth used to generate the observed image. In the present example, one particle was reconstructed with a false diameter but correct position, and is marked by the red circle."[SGK25]. Figure from [SGK25]	92
4.18	"Example of the use of the IP approach on real experimental data. The observed image is from the bubble experiment. The observed image is fitted by minimizing the mismatch functional \mathcal{P} . The FM is used to generate an image with the optimal particle parameter set ψ_I . This image can be used as a reference to qualitatively check the result of the IP approach. The particle positions and diameters are then directly obtained from ψ_I ."[SGK25]. Figure from [SGK25]	92
5.1	"Three different types of PIs: PIs with an empty pattern (left), PIs with a regular unidirectional pattern (middle) and PIs with an irregular speckle pattern (right). The particle type, the scattering characteristic and simulated and real examples are shown. The simulated examples were created using the forward model from Sax <i>et al.</i> [SGK25] and the glare point based scattering characteristic. For the rough opaque particle a number of glare points significantly larger than two (here 12) were chosen to model multiple randomly distributed glare points. The real examples show a DEHS tracer from Exp. G2 (left), a water droplet from Exp. T4 (middle) and a Polyamid tracer from Exp. G1 (right)."[SK25b]. Figure from [SK25b].	98
5.2	"Auto-labeling approach and PI detection. Two data acquisition experiments are conducted, each involving a single phase. PIs are extracted from the resulting image sets using conventional algorithms. These extracted PIs are then used to train GANs, to generate mixed-phase image sets. The mixed-phase sets are subsequently used to train an object detection CNN. Finally, the trained CNN is applied to the actual two-phase experimental data."[SK25b]. Figure from [SK25b]	99
5.3	"Generation of automatically labeled training data for object detection networks. Two GANs - one for tracer particles and one for the dispersed phase - independently generate small image snippets, each containing a single PI. For each snippet, a BB is placed over the PI. The image snippets and corresponding BBs are then randomly resized and inserted into an empty image."[SK25b]. Figure from [SK25b].	100
5.4	"PIs sampled from the Exp. G1[LSBK24] (a), Exp. G2[149] (b), Exp. G3[SDK25] (c) and Exp. G4[LSBK24] (d). These PIs were used to train the GANs for tracer particles (a)&(b) and for particles of the dispersed phase (c)&(d)."[SK25b]. Figure from [SK25b].	102
5.5	"GAN-generated images of PIs of tracer particles (a) and particles of the dispersed phase (b). All images have dimensions of 64×64 px."[SK25b]. Figure from [SK25b].	103
5.6	"Exemplary images from the six test datasets: T1[LSBK24] (a) and T2[115] (b) show tracer particles in a single phase flow. T3[SDK25] (c) and T4[151] (d) show only bubbles and droplets but no tracers. T5 (e) shows a synthetic image of a two phase flow of bubbles and tracers and T6[LSBK24] (f) show a real image of bubbles and tracers."[SK25b]. Figure from [SK25b].	105

6.1	"Three issues caused by PI overlap: The detection of covered PIs is more difficult (left). The PI size measurement becomes more challenging and the PI size might be only estimated due to the PI boundaries being covered (middle). The free remaining area (FRA) that can be used in IPI becomes small due to PI overlap (right)."[SK25a]. Figure from [SK25a].	114
6.2	"Overview over the dataset for the empirical investigation of the PI overlap visualized. Six different number of particles per image $N_P = [10, 20, 30, 40, 50, 60]$ and ten different defocus lengths (i.e. PI sizes) $z_P = [1, 1.5, 2, 3, 4, 5, 6, 7, 8, 9]$ mm were investigated, resulting in 6×10 sets of 400 images each. The images were generated using the physical model from Sax <i>et al.</i> [SGK25] with parameters from the experiment described in [SDK25]."[SK25a]. Modified figure from [SK25a].	115
6.3	"Exemplary images from the four different experiments used to validate the PI overlap model. , Exp. 1[115] shows tracers in a gap flow (a), Exp. 2[LSBK24] shows tracers and bubbles in a vortex like rotating flow (b) and Exp. 3[SDK25] shows shows a bubble column in a resting fluid (c), Exp. 4[151] shows droplets in a spray behind a nozzle (d)."[SK25a]. Figure from [SK25a].	116
6.4	"Different measures to describe first degree PI overlaps: IoU (left), IoM (middle) and IoA (right). The issues with the IoU and the IoM are visualized."[SK25a]. Modified figure from [SK25a].	118
6.5	"First and higher degree PI overlaps visualized (left). Metrics for higher degree PI overlaps are visualized (right). The FRA and TCA are shown on the example of a PI in IPI."[SK25a]. Modified figure from [SK25a].	120
6.6	"A higher degree PI overlap visualized. The gray area is the PI for which the PI overlap is regarded. The red area is the intersection area, and the blue area shows the twice overlapped area of the PI. The summed IoA overestimated the TCA of the PI by the blue area."[SK25a]. Modified figure from [SK25a].	120
6.7	"Deterministic pixel counting approach by simulating discretized circles. A circle consists of pixels with a value of one (marked as gray). The overlap areas of each circle are obtained from the number of pixels with a value of larger than one (marked as red). Different grid resolutions of the simulation are shown from 100 px to 400 px."[SK25a]. Figure from [SK25a].	121
6.8	"The approximation of the TCA by the quadrature-like method for $N_P = 60$. The discretized simulations are shown for different grid resolutions $N_{px} = \{200, 300, 400\}$. It can be seen that almost no difference between the $N_{px} = 300$ and $N_{px} = 400$ resolution can be observed."[SK25a]. Figure from [SK25a].	122
6.9	"The mean number of overlaps per PI over the number of PIs for different defocus distances (a) and over the defocus distance for different number of PIs (b)."[SK25a]. Figure from [SK25a].	124
6.10	"The mean number of overlaps per PI over the seeding density $\mathcal{S} \sim N_P z_P^2$ for different number of PIs in the image (a). The mean fraction of PIs without overlap in an image over the seeding density \mathcal{S} and for different number of PIs in the image (b). Shown are also the fits to the curves (<i>dashed line</i> ---), which are specified in Tab.6.1. The values obtained from the validation experiments are provided by the markers ($\triangle, \triangleleft, \diamond, \triangleright$). Images from the validation experiments are shown in Fig. 6.3."[SK25a]. Figure from [SK25a].	124
6.11	"The max IoA per PI over the seeding density $\mathcal{S} \sim N_P z_P^2$ for different number of PIs in the image (a). The mean TCA with $N_{px} = 400$ over the seeding density \mathcal{S} for different number of PIs in the image (b). Shown is also the fit to the curves (<i>dashed line</i> ---), which is specified in Tab.6.1. The values obtained from the validation experiments are provided by the markers ($\triangle, \triangleleft, \diamond, \triangleright$). Images from the validation experiments are shown in Fig. 6.3."[SK25a]. Figure from [SK25a].	125

7.1	Overview of the measurement framework for DTPFs in optically constrained environments. A CNN is used for the initial detection and phase separation. The BBs outputted by the CNN are used to initialise the IP approach. Once the IP approach converges, the final particle positions and sizes are derived from the updated input vector. The contributions of each chapter are indicated in grey within the schematic. Modified figure adapted from [SGK25].	131
B.1	"Generator score (G-score) and discriminator score (D-score) during the training of the two GANs. Two GANs were trained: One for the generation of PIs representing tracers (a) and another to represent particles of the dispersed phase (b). In both cases the D-Score converged to values close to 0.5 indicating that the generator successfully fools the discriminator." [SK25b]. Figure from [SK25b].	182
B.2	"Total training loss of Faster R-CNN (a) and YOLOv4 and tiny YOLOv4 (b)." [SK25b]. Figure from [SK25b].	183
B.3	"False positive class biases for Exps. T1, T2, T3, T4, T5, T6 over the applied confidence score threshold. The black dashed line indicates zero bias. The bias becomes +1 if all FP are predicted to be tracers and -1 if predicted as dispersed phase." [SK25b]. Figure from [SK25b].	184
C.1	"Test for statistical convergence by investigating the mean over the standard deviation of the overlap fraction for increasing dataset size. The x -axis is showing the number of PIs in the dataset. The depicted overlap fraction was computed by the discretized direct computation with an image of resolution 400 px. The test were conducted for $N_P = 10, 20, 30, 40, 50, 60$ particles per image (a),(b),(c),(d),(e) and (f) respectively, and for ten different PI sizes." [SK25a]. Figure from [SK25a]. . .	185

List of Tables

3.1	"Uncertainty as a systematic error introduced by an angle uncertainty of 1° "[SDK25]. Table from [SDK25]	61
4.1	"Median Error for 200 synthetic PIs with random particle parameters ψ_P . The particles have diameters in the range of $d_P = [20, 300] \mu\text{m}$ and a random distance from the focal plane in the range of $z_P = [1, 8] \text{mm}$. For the FFT approach a random error within an uncertainty of $\Delta x_{PI}, \Delta y_{PI}, \Delta d_{PI} = 1 \text{px}$ was added to the ground truth parameters ψ_P . The FFT was then performed on the initialised value. For the IP approach a random placement error for the BB $\Delta x_{PI}, \Delta y_{PI}, \Delta d_{PI} = 1 \text{px}$ and 5px was used to initialise the starting values of the optimization. The factor of improvement for the IP based on the chosen FFT+BB (with 1px error) approach is given in the last two lines of the table. The improvement factor of the IP with 5px initialization error relates to an improvement over the FFT+BB with 1px error."[SGK25]. Table from [SGK25]	84
4.2	"Median errors of the FFT + BB (1px uncertainty) and the IP approach (initialisation with 1px and 5px uncertainty). The factor of improvement of the result using the IP approach is given below. The number of outlier is stated, with an outlier, arbitrarily being defined as an error in diameter of more than $10 \mu\text{m}$. The outlier threshold of $10 \mu\text{m}$ was chosen based on the PDFs in Fig. 4.14c, where two ranges of high probability are separated: one in the order of magnitude of $1 \mu\text{m}$ and the other in the order of magnitude of $100 \mu\text{m}$. The improvement factor of the IP with 5px initialization error relates to an improvement over the FFT+BB with 1px error."[SGK25]. Figure from [SGK25]	89
5.1	"Class agnostic PI detection performance for all six test sets. Shown are the average precision, the max recall and the normalized truncated average precision. The best score for each metric and each test set is highlighted in bold print."[SK25b]. Table from [SK25b].	108
5.2	"Classification results on the six different test datasets. The class accuracy for true positive detections and the class bias of false positives is given. The bias becomes +1 if all FP are predicted to be tracers and -1 if predicted as dispersed phase. The best score for each metric and each test set are highlighted in bold print."[SK25b]. Table from [SK25b].	109
6.1	"The different metrics to describe overlap and their respective models are shown. The model type as well as the fit function and the parameters are given. The goodness of fit is measured by R-squared."[SK25a]. Table from [SK25a].	126
A.1	k and q values for the calculation of $\beta_i^{(p)}$ and $\beta_t^{(p)}$ in Eq. 2.13. Table from [SDK25] . . .	178
B.1	"Specifications of the generator (a) and discriminator (b) network in the GAN, used for both the tracer GAN and the GAN for the dispersed phase. For more information see [131]."[SK25b]. Table from [SK25b].	181
B.2	"Training Specifications for the GAN network. For more information see [131]."[SK25b]. Table from [SK25b]	182

B.3	"Training settings for Faster R-CNN (ResNet50), Tiny YOLOv4 and YOLOv4 (DarkNet)" [SK25b]. Table from [SK25b].	183
-----	--	-----

A Appendix: Interferometric Particle Imaging in Front-, Side- and Backscatter

The content of this appendix is based on the article *Interferometric Particle Imaging for Particle Sizing in the Front-, Side-, and Backscatter Region - Towards Single Optical Access IPI* by Sax *et al.* [SDK25].

A.1 Further Information on the Optical System of the Validation Experiment

"Eq. (3.4) requires a description of the optical system (B_{tot} - defocusing component) to calculate the particle diameter. This section provides more details on the computation of the value of B_{tot} in the validation experiment. Three different kinds of matrices were used. The first one is a matrix for the propagation through free space of length z_i .

$$M_{\text{p},i} = \begin{bmatrix} 1 & z_i \\ 0 & 1 \end{bmatrix} \quad (\text{A.1})$$

The next matrix describes the refraction at an interface like a glass wall. The initial refractive index is n_1 and the index after the interface is n_2 . For a curved wall the radius $R > 0$ is used to describe the curvature. For a flat interface $R \rightarrow \infty$ is used.

$$M_{\text{I}1 \rightarrow 2} = \begin{bmatrix} 1 & 0 \\ \frac{n_1 - n_2}{R_1 n_2} & \frac{n_1}{n_2} \end{bmatrix} \quad (\text{A.2})$$

The third required matrix describes a thin lens with focal length f_{lens} .

$$M_{\text{lens}} = \begin{bmatrix} 1 & 0 \\ -\frac{1}{f_{\text{lens}}} & 1 \end{bmatrix} \quad (\text{A.3})$$

The optical system comprises of propagation through free space from the bubble to the glass wall ($M_{\text{p},\text{w}}$), the interface between the water and the glass ($M_{\text{I},\text{w} \rightarrow \text{g}}$), the propagation through the glass ($M_{\text{p},\text{g}}$), the interface between glass and air ($M_{\text{I},\text{g} \rightarrow \text{a}}$), propagation from the glass wall to the imaging lens ($M_{\text{p},\text{a}}$), the changing of the light path at the imaging lens (M_{lens}) and finally the propagation from the lens onto the camera chip ($M_{\text{p},\text{im}}$), see Fig. 3.9c. The total system is then computed from

$$M_{\text{tot}} = M_{\text{p},\text{im}} M_{\text{lens}} M_{\text{p},\text{a}} M_{\text{I},\text{g} \rightarrow \text{a}} M_{\text{p},\text{g}} M_{\text{I},\text{w} \rightarrow \text{g}} M_{\text{p},\text{w}}. \quad (\text{A.4})$$

"[SDK25]

A.2 Algorithm for the Calculation of Glare Point Maps

"The calculation of the glare point maps (GP-maps) as conducted alongside Eq. (2.13) is based on

$$\theta' = 2(\beta_i^{(p)} - p\beta_t^{(p)}) = 2\pi k + q\theta \quad (\text{A.5})$$

and Snell's law. As mentioned before, q is ± 1 and $k \in \mathbb{N}$ is an integer. The values of q and k remain to be determined, where only one combination of q and k satisfies the equation for a given refractive index m and scattering order p . A method to determine q and k while computing $w(\theta)$ is demonstrated in Algorithm 1, which is based on a selective trial and error approach.

Algorithm 1 work as follows: For a given scattering order p and refractive index m , Eq. (2.13) is solved indirectly, by inserting values of $\beta_i^{(p)} = [0, \pi/2]$ into the equation and calculating the resulting values of θ . This is done for every value of $q = \{-1, 1\}$, $k = [0, 1, 2, \dots, k_N]$ (for low scattering orders $k_N = 5$ is sufficient) in a trial and error approach. This results in $N_{\text{tries}} = 2k_N$ different arrays for $\theta \left(\beta_i^{(p)} \right)$. The correct values of q and k can be determined by constraints on θ . The constraints for correct solutions of the equation are $\Re\{\theta\} \in [0, \pi]$ and $\Im\{\theta\} = 0$, which is only satisfied for one pair of values of q and k .

Also, not all values of $\beta_i^{(p)} = [0, \pi/2]$ result in $\Re\{\theta\} \in [0, \pi]$ and $\Im\{\theta\} = 0$, since the required conditions for a $\beta_i^{(p)}$ -determination are not given for the full range $[0, \pi/2]$. The correct entries of the $\beta_i^{(p)}$ array can be identified by the entries that satisfy the constraints on θ . The truncated arrays θ_{output} and β_{output} are defined as the set of values in θ and β that satisfy the constraints on θ . The glare point position w of order p is then computed by inserting the values of β_{output} and q into Eq. (2.9).

Algorithm 1 returns the two arrays of θ_{output} and w_p , and the used values of q and k . The values of q and k for $m > 1$ and $m < 1$ are given for the first four scattering orders $p = [0, 3]$ in Tab A.1."[SDK25]

Table A.1: k and q values for the calculation of $\beta_i^{(p)}$ and $\beta_t^{(p)}$ in Eq. 2.13. Table from [SDK25]

p	k for $(m > 1)$	q for $(m > 1)$	k for $(m < 1)$	q for $(m < 1)$
0	0	1	0	1
1	0	-1	0	1
2	0	-1	0	-1
3	-1	1	1	1

Algorithm 1 Calculation of a line with scattering order p in a GP-map. Algorithm from [SDK25]

Input: $p = p_i$ # defines the scattering order
 $q = [-1, 1]$
 $k = [0, 1, \dots, k_N]$ # k_N is the number of tries for k
 $\beta_i = \text{linspace}(0 \dots \pi/2, N_{\text{setps}})$ # N_{setps} defines the resolution of the GP-map

- 1 **for** i *in* q **do**
- 2 **for** j *in* k **do**
- 3 $\theta(:, i, j) = q(i)/2 \cdot (\beta_i(:) - p \cdot \arccos(1/m \cos(\beta_i(:)) - \pi k(j))$
- 4 $[i_\beta(:), i_q, j_k] = \text{find}(\Re\{\theta(:, i, j)\} \geq 0 \ \& \ \Re\{\theta(:, i, j)\} \leq \pi \ \& \ \Im\{\theta(:, i, j)\} = 0)$
get entries in β_i , k and q arrays
 $q_{\text{output}} = q(i_q)$ # get correct q
 $k_{\text{output}} = k(j_k)$ # get correct k
 $\theta_{\text{output}}(:) = \theta(i_\beta(:), i_q, j_k)$ # get scattering angles that match w_p
 $w_p(:) = q(i_q) \cdot \cos(\beta_i(i_\beta(:)))$ # get glare point positions
- 5 **return** $q_{\text{output}}, k_{\text{output}}, \theta_{\text{output}}, w_p$

B Appendix: Neural Networks and Training for Two-Phase DPTV

The content of this appendix is based on the article *Pattern-Based Phase-Separation of Tracer and Dispersed Phase Particles in Two-Phase Defocusing Particle Tracking Velocimetry* by Sax *et al.* [SK25b].

B.1 GAN Training Configuration and Training Scores

"This section provides further details on the GAN architecture (Tab. B.1) and the GAN training process (Tab. B.2 and Fig. B.1). The generator network consists of a project and reshape layer that transforms the latent noise vector, followed by three blocks of transposed convolutional layers, each with batch normalization and ReLU activation. The final layer is a transposed convolution with a `tanh` activation function. The discriminator network begins with a dropout layer at the input, followed by four convolutional layers with batch normalization and ReLU activation. The final layer is a convolutional layer with a `sigmoid` activation function. The generator and discriminator architectures used in the present work follow the example provided by MATLAB [131]."[SK25b]


Table B.1: "Specifications of the generator (a) and discriminator (b) network in the GAN, used for both the tracer GAN and the GAN for the dispersed phase. For more information see [131]."[SK25b]. Table from [SK25b].

Network Architecture	Project and Reshape 3× (Transposed Conv. Batch Normalization ReLU) Transposed Conv. tanh	Network Architecture	Dropout Convolution ReLU 2× (Convolution Batch Normalization ReLU) Convolution Sigmoid
Dim Latent Input	100	Input Size	[64,64,1]
Output Size	[64,64,1]	Output	probability
Filter Size (Conv)	5	Dropout Probability	0.5
Num Filters (Conv)	64	nNmFilters (Conv)	64
Stride (Conv)	2	FilterSize (Conv)	5
		Stride (Conv)	2
		Scale (Leaky ReLu)	0.2
(a) generator		(b) discriminator	

"The training parameters for the GANs are provided in Tab. B.2 for both the tracer GAN and the dispersed phase GAN. The GANs were trained using the ADAM optimizer [95]. The generator and discriminator scores during training are shown in Fig. B.1."[SK25b]

Table B.2: "Training Specifications for the GAN network. For more information see [131]."[SK25b]. Table from [SK25b]

Network	GAN (Tracer)	GAN (Dispersed Phase)
solver	ADAM	ADAM
miniBatchSize	128	128
learnRate	2.0000e-04	2.0000e-04
gradientDecayFactor	0.5000	0.5000
squaredGradientDecayFactor	0.9990	0.9990
labels flipProb	0.3500	0.3500
max Epoch	50	50
Output Epoch	50	28
Output Iteration	11150	6756
data augmentation	horizontal flip vertical flip	horizontal flip vertical flip


**Figure B.1:** "Generator score (G-score) and discriminator score (D-score) during the training of the two GANs. Two GANs were trained: One for the generation of PIs representing tracers (a) and another to represent particles of the dispersed phase (b). In both cases the D-Score converged to values close to 0.5 indicating that the generator successfully fools the discriminator."[SK25b]. Figure from [SK25b].

B.2 Training Losses and Configuration for the Object Detection Networks

"The training configuration and the hyperparameters used to train Faster R-CNN and YOLOv4 are presented in Tab. B.3. Both networks were trained using stochastic gradient descent with momentum. Pre-trained weights from MS-COCO 2018 were used to initialize the training. The corresponding training losses are shown in Fig. B.2."[SK25b]

Table B.3: "Training settings for Faster R-CNN (ResNet50), Tiny YOLOv4 and YOLOv4 (DarkNet)"[SK25b]. Table from [SK25b].

$IoU \geq 0.5$	Faster R-CNN	YOLOv4	YOLOv4
	ResNet50	CSPDarkNet53-Tiny	CSPDarkNet53
parameters (millions)	42	5.8	63.9
optimizer	SGDM	SGDM	SGDM
training iterations	20,000	20,000	20,000
patience (validation checks)	none	50	50
validation frequency	50	50	50
output iteration	20,000	10,850	20,000
minibatch size	4	4	4
learning rate α	0.0003	0.001	0.001
learning rate decay	none	none	none
momentum β	0.9	0.9	0.9
L_2 -regularization	none	0.0005	0.0005
data augmentation	horizontal flip vertical flip	horizontal flip vertical flip	horizontal flip vertical flip
pre-trained	MS-COCO2018	MS-COCO2018	MS-COCO2018

**Figure B.2:** "Total training loss of Faster R-CNN (a) and YOLOv4 and tiny YOLOv4 (b)."[SK25b]. Figure from [SK25b].

B.3 False Positive Class Bias

"For the three object detection networks, Faster R-CNN, YOLOv4, and Tiny YOLOv4, the false positive class bias was computed and is presented in Fig.B.3. It can be observed that in Exps. 3, 4, 5, and 6, the class bias starts near zero for confidence scores of 0.0001 and approaches -1 as the confidence score increases to 0.999. This indicates that the networks exhibit increasing confidence in the dispersed phase and are biased toward this class. In the single-phase experiments containing only tracers (Exps. 1 and 2), a class bias toward the tracer class is observed."[SK25b]




Figure B.3: "False positive class biases for Exps. T1, T2, T3, T4, T5, T6 over the applied confidence score threshold. The black dashed line indicates zero bias. The bias becomes +1 if all FP are predicted to be tracers and -1 if predicted as dispersed phase." [SK25b]. Figure from [SK25b].

C Appendix: Particle Image Overlap in Single Camera Defocusing Approaches without Astigmatism

The content of this chapter is based on the article *On the Particle Image Overlap in Single Camera Defocusing Approaches* by Sax *et al.* [SK25a].

"The test for statistical convergence of the overlap fraction is shown in Fig. C.1. The mean normalized by the standard deviation of the total covered area (TCA) is plotted for increased dataset size. It can be seen that a dataset size of 400 images with $N_P = \{10, 20, 30, 40, 50, 60\}$ PIs per image is sufficiently large to observe no changes in the TCA with further increasing dataset size." [SK25a]




Figure C.1: "Test for statistical convergence by investigating the mean over the standard deviation of the overlap fraction for increasing dataset size. The x -axis is showing the number of PIs in the dataset. The depicted overlap fraction was computed by the discretized direct computation with an image of resolution 400 px. The test were conducted for $N_P = 10, 20, 30, 40, 50, 60$ particles per image (a),(b),(c),(d),(e) and (f) respectively, and for ten different PI sizes." [SK25a]. Figure from [SK25a].