ETH Zürich
Institute of Fluid Dynamics

Intercomparison of Snow Imaging and Water Jet Tracking Experiments

Jules Cavin



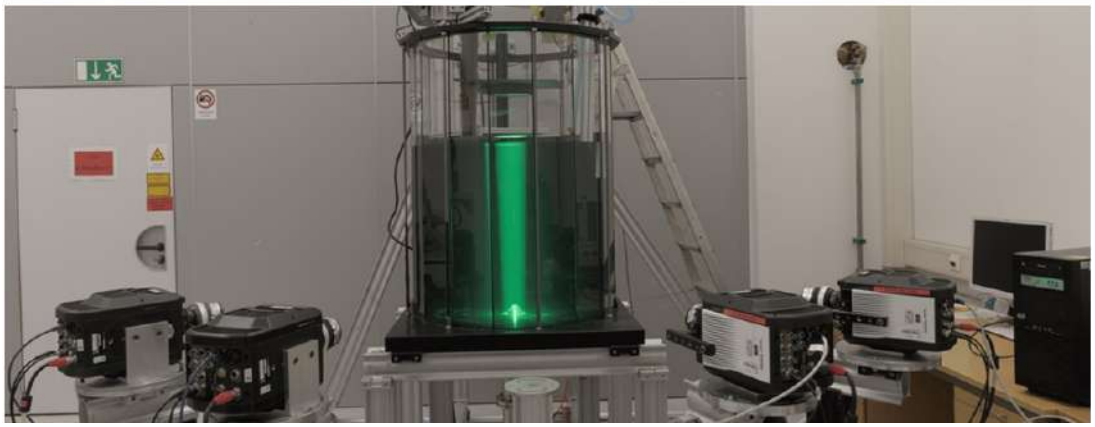

Bachelor Thesis

Professor: Prof. Dr. Filippo Coletti
Supervisor ETH Zürich: Dr. Koen Muller
Supervisor DLR Göttingen: Prof. Dr. Andreas Schröder;
Dr. Daniel Schanz

Spring Semester 2026

Acknowledgements

I would like to express my gratitude to my supervisor, Dr. Koen Muller, for his highly valuable guidance, support, and patience throughout this bachelor's thesis. His expertise and feedback were crucial in shaping both the experimental work and my understanding of the underlying concepts, allowing me to work under ideal conditions throughout this project.

I'm also deeply grateful to Prof. Dr. Andreas Schröder and Dr. Daniel Schanz for welcoming me to DLR Göttingen and allowing me to carry out part of my thesis work there. Their expertise and support during my stay were essential for the completion of this work.

A special thanks goes to Jonathan Lemarechal, who was the first to welcome me to DLR Göttingen. He introduced me to the laboratory, guided me through the experimental setup, and worked alongside me during the water-jet experiments. His patience, enthusiasm, and willingness to share his knowledge made my stay at DLR both enjoyable and highly educational.

Conducting the field experiments during cold winter nights in the Swiss Alps was no easy task. I would therefore like to sincerely thank Ramon Carabelli, Doruk Emür, Alessandro Gambino, Sergi González-Herrero, Ivan Calic, and Udhav Ulhas Gawandalkar for their invaluable support, teamwork, and companionship throughout these field campaigns.

I would also like to sincerely thank Prof. Dr. Filippo Coletti for making this project possible and for giving me the opportunity to contribute to this exciting research project.

Finally, I would like to express my heartfelt gratitude to all the friends and wonderful people who accompanied me throughout this journey. Their friendship, support, and the countless memorable moments we shared made this experience truly unforgettable. In particular, I would like to thank Flora Könz for her constant encouragement, unwavering support, and invaluable help in designing some of the figures presented in this thesis.

Abstract

This thesis compares two applications of volumetric Lagrangian particle tracking: natural snowfall in the atmospheric boundary layer and a turbulent round water jet under controlled laboratory conditions. The aim is to assess how calibration accuracy, imaging geometry, particle distribution, and reconstruction artefacts differ between a large-scale outdoor measurement and a laboratory benchmark, with particular attention to the feasibility of applying Shake-The-Box (STB) techniques to natural snowfall.

The snowfall experiment used sixteen cameras arranged in four arrays to image a 10 $\times$ 10 $\times$ 10 m^3 measurement volume in Davos, Switzerland. A drone-borne tetrahedral light-marker system provided the calibration data, and an in-house Lagrangian tracking framework was used to reconstruct snowflake trajectories. The laboratory experiment, performed at DLR Göttingen, tracked 50 µm tracer particles in a turbulent round water jet over a cylinder with a diameter of 8 centimetres and a height of 30 centimetres at Reynolds numbers of 2000, 4000, and 11000. The water-jet data were processed with a target-based calibration refined by volume self-calibration and optical-transfer-function reconstruction, followed by STB processing in DaVis.

The laboratory setup achieved reprojection errors of only a few hundredths of a pixel and reproduced the expected centerline decay and self-similar radial profiles of a turbulent round jet, confirming it as a reliable benchmark. In the snowfall experiment, reprojection errors were reduced from roughly four to five pixels to about two to three pixels using Soloff polynomial mappings, corresponding to sub-snowflake accuracy. However, the outdoor reconstruction remained affected by non-uniform particle coverage, line-of-sight ambiguities, and likely ghost-particle contamination, especially because only short four-frame tracks and a reduced set of camera views were available. Overall, the results show that large-volume snow tracking can reach promising calibration accuracy, but further work, particularly optical-transfer-function reconstruction and a first STB processing of snow data, is required before fully quantitative high-density snowfall tracking can be achieved.

Acronyms

ABL Atmospheric Boundary Layer
BLE Boundary Layer Equations
DLR German Aerospace Center
FF3 FlowFit3
FOV Field Of View
FSR Fine Scale Reconstruction
IFD Institute of Fluid Dynamics
IPR Iterative Particle Reconstruction
LPT Lagrangian Particle Tracking
OTF Optical Transfer Function
PDF Probability Density Function
ppp Particles-per-pixel
PTV Particle Tracking Velocimetry
STB Shake The Box
VSC Volume Self Calibration

List of Symbols

d	Jet nozzle's diameter
f	Focal length
$K^\star$	Intrinsic camera calibration matrix
ℓ_m^+	Mixing length
P	Extrinsic projection matrix
r	Radial coordinate
R	Rotation matrix
$r_{1/2}$	Jet's half-width
s	Camera skew
S	Spreading rate
$\mathbf{t}$	Translation vector
$\mathbf{T}$	Position vector
u_τ	Friction velocity
u^+	Dimensionless velocity
$\langle U \rangle$	Mean axial velocity
U_J	Axial velocity at the jet exit
U_0	Centerline velocity
U_∞	Free-stream velocity
x	Axial coordinate
$\boldsymbol{x}$	Coordinate in camera frame
$\tilde{\boldsymbol{x}}$	Coordinate in camera frame in augmented form
x_0	Virtual origin
$\boldsymbol{X}$	Coordinates in object space
y^+	Dimensionless wall distance
$\alpha_{x,y}$	Pixel pitch
δ	Boundary layer thickness
θ	Azimuthal coordinate
κ	Von Kármán constant
ν	Fluid kinematic viscosity
ρ	Fluid density
τ_ω	Wall shear stress

Contents

1 Introduction **1**

1.1 The Atmospheric Boundary Layer 2
1.2 The Free Shear Turbulent Jet Flow 4
1.3 Particle Tracking Velocimetry 6
1.4 Shake the Box 6
1.5 FlowFit 8
1.6 Goal and Outline 9

2 In the Field and in the Lab **10**

2.1 Outdoor Snow Tracking Experiment 10
2.1.1 Experimental Layout 11
2.1.2 Drone-based Camera Calibration 12
2.1.3 Snowfall Measurement Campaign 12
2.2 Laboratory Water Jet Tracking Experiment 13
2.2.1 Experimental Layout 13
2.2.2 Illumination and Optical System 15
2.2.3 Camera Arrangement and Calibration 15
2.2.4 Water Jet Measurement Campaign 16

3 Data Processing and Calibration **17**

3.1 Snow Tracking 17
3.1.1 Drone Calibration Framework 17
3.1.2 Snow Tracking Reconstruction 20
3.2 Turbulent Water Jet 20
3.2.1 Calibration using Target 20
3.2.2 Volume Self Calibration 21
3.2.3 Optical Transfer Function 22
3.2.4 Commercial STB DaVis 23

4 Results **26**

4.1 Drone Calibration Results 26

4.1.1 Object Skew Error . . . 26
4.1.2 Reprojection Error . . . 27
4.1.3 Soloff's Reprojection Error . . . 28
4.2 Snow Tracking Results . . . 28
4.2.1 Snowflakes Reprojection . . . 30
4.2.2 Snow Tracking Data . . . 32
4.2.3 Extracting Field Descriptions . . . 35
4.3 Water Jet Tracking Results . . . 37
4.3.1 STB Data from DaVis . . . 37
4.3.2 Fine Scale Reconstruction Data . . . 40

5 Discussion **42**
5.1 Comparison between Field and Lab Experiments . . . 42
5.1.1 Calibration Accuracy and Measurement Uncertainty . . . 42
5.1.2 Spatial Distribution of Particles . . . 43
5.1.3 Line of Sight Ambiguities and Occlusions . . . 45
5.1.4 Ghost Particles in the Water Jet Data . . . 46
5.2 Analysis of the Flows Physics . . . 47
5.2.1 Snow Boundary Layer . . . 48
5.2.2 Water Jet Velocity Profiles . . . 49

6 Conclusion **51**
6.1 Snow Tracking STB Feasibility . . . 53

References **53**

Declaration of Originality **62**

Chapter 1

Introduction

Falling snowflakes in the atmosphere reveal the physical behaviour of turbulence, with applications ranging from heat transfer and pollutant dispersion to weather forecasting [1]. Understanding the complex dynamics of falling snowflakes in turbulent air can be of significant interest, as the latter cannot be captured by numerical models without substantial simplifications [2]. Field experiments offer a novel solution to track the motion of natural snowflakes in the atmospheric boundary layer and have gained increasing attention as a complementary approach to laboratory experiments and numerical modelling. Experiments allow us to directly study interactions between particles with different settling velocities, which, due to their differing inertias, influence snow deposition patterns and clustering. This is crucial for understanding and anticipating natural hazards such as avalanches [3]. In mountainous regions, avalanches represent a major risk to human life, infrastructure, and transportation networks, and their release is strongly dependent on how the snow is deposited and redistributed by the wind [4]. Improving our knowledge of snowflake behaviour in the atmosphere is therefore essential to better anticipate these phenomena and reduce accidents and risks.

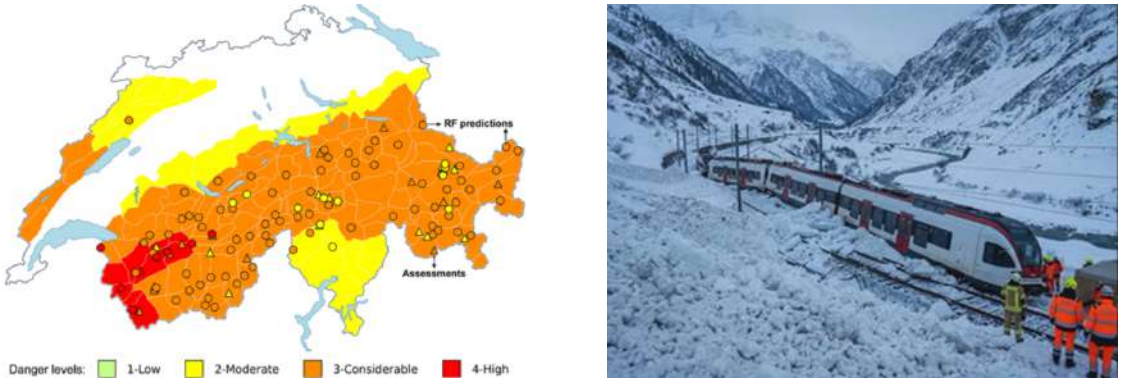

(a) (b)

Figure 1.1: (a) Avalanche risk in Switzerland on 6 March 2020. Adopted from [5]. (b) A train derailed due to an avalanche on 16 February 2026 in Goppenstein, Switzerland.

The first optical field measurements of snow settling in the atmosphere were performed in 2017 and reported by [6], in which snowflakes were imaged over a vertical plane of approximately 30 m^2, demonstrating, for the first time in a natural setting, that snowflakes behave as inertial particles in turbulence. Three field campaigns conducted in [7] provided, for the first time, quantitative evidence of snow particle clustering in the atmosphere and further advanced the understanding of snowflake dynamics in atmospheric turbulence.

Both these studies, however, were limited to two-dimensional measurements on vertical planes. This limitation was addressed by [8] in 2023, who developed the first three-dimensional particle tracking system capable of reconstructing snow trajectories within a volume of approximately 4 x 4 x 6 m^3 allowing snow tracking over a further restricted volume of 8 m^3. The present work aims to extend this line of investigation by significantly increasing the three-dimensional measurement volume, tracking snowflakes over a 10 x 10 x 10 m^3 domain.

1.1 The Atmospheric Boundary Layer

Tracking snow particles outdoors represents a general measurement application in the Atmospheric Boundary Layer (ABL). To place these observations into a clear physical framework, we first introduce its basic structure and dynamics.

The ABL, or planetary boundary layer, is the lowest part of the troposphere that is in direct contact with the Earth's surface, where surface heating, cooling, and shear generate and sustain turbulent motions [4].

Above the boundary layer lies a stable free atmosphere, where turbulence from below generally cannot penetrate far; the height of this region is very variable, depending on temperature, weather conditions, and obstacles on the ground. It ranges from a few hundred metres to a few kilometres [9].

From a physics perspective, the simplest boundary layer model is one where a uniform-velocity U_∞ non-turbulent stream flows over a flat surface. In this case, we are interested in the vertical profile of the horizontal flow velocity u within the boundary layer, which can be decomposed into three regions: the viscous sublayer close to the wall, an intermediate buffer layer, and an outer region that comprises the logarithmic (log-law) layer and the outer layer of the turbulent boundary layer [10]. All these expressions can be regarded as approximate solutions of the Boundary Layer Equations (BLE) under different assumptions [11, p. 113].

Viscous sublayer The viscous sublayer occurs at $y^+ < 5$, where y^+ is the dimensionless wall distance, defined as $y^+ = yu_\tau/\nu$, with y being the distance from the ground, ν the

kinematic viscosity, and $u_\tau = \sqrt{\tau_\omega/\rho}$ the friction velocity. In this region, viscosity is dominant, and the equation for the scaled velocity profile $u^+ = u/u_\tau$ is

$$u^+ \approx y^+ \tag{1.1}$$

Logarithmic layer The log law holds for $y^+ > 50$ and $y/\delta \ll 1$, where δ is the boundary layer thickness, and has the following form:

$$u^+ = \frac{1}{\kappa}\ln(y^+) + B \tag{1.2}$$

Here κ is the von Kármán constant (≈ 0.41), and B is an integration constant. For flat surfaces, it is $B \approx 5.0 - 5.5$

Buffer layer At last, we are also interested in the buffer layer ($5 < y^+ < 50$); an empirical answer for the velocity in this range is given by von Driest [12]:

$$u^+ = \int_0^{y^+} \frac{2\mathrm{d}y'}{1 + [1 + 4\ell_m^+(y')^2]^{\frac{1}{2}}} \tag{1.3}$$

Where the mixing length is $\ell_m^+ = \kappa y^+[1 - \exp(-y^+/A^+)]$, with the constant $A^+ = 26$. It can be seen that for large y^+ the negative exponential function goes to zero, and the log law is recovered, with $B = 5.3$ [11].

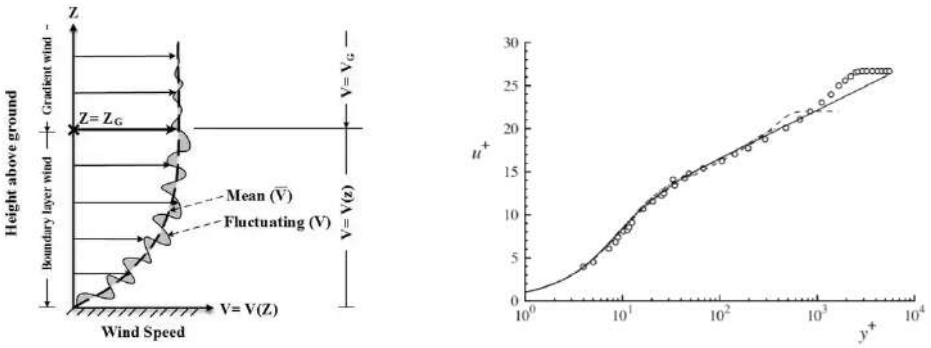

(a) (b)

Figure 1.2: (a) Schematic of the vertical wind speed profile in the ABL, showing the increase from the surface within the turbulent boundary layer and the quasi-uniform flow above. Reproduced from [13]. (b) Dimensionless mean velocity u^+ as a function of the dimensionless wall distance y^+. The circles, dashed line, and dot–dashed line represent experimental data at different Reynolds numbers, while the solid line shows von Driest's law. Reproduced from [11]

1.2 The Free Shear Turbulent Jet Flow

To compare snow tracking experiments in the ABL with the performance and accuracy of laboratory measurements, this thesis also considers a turbulent round jet, a canonical free-shear flow configuration that is well studied in fluid mechanics and yields closed-form solutions for its mean velocity profile [11].

Turbulent jet flows are generated by a jet of fluid emerging from an opening with a velocity significantly higher than that of the surrounding fluid, resulting in turbulent motions at the interface between the two regions. The name "free" implies that the jet flow is remote from walls, such that turbulence arises only from the mean-velocity differences. The study of this particular flow configuration offers several engineering benefits, ranging from an improved understanding of mixing and heat transfer in industrial devices to a convenient physical model of rotor-induced flow [14].

A key feature of turbulent round jets is the axial evolution of the centerline velocity $U_0(x)$; experimental data show that the inverse of the centerline velocity varies linearly with the axial distance (see Figure 1.3), leading to the following relation:

$$\frac{U_0(x)}{U_J} = \frac{B}{(x - x_0)/d} \tag{1.4}$$

Where U_J is the axial velocity at the jet exit, d is the nozzle diameter, B is an empirical constant, and x_0 is the intersection of the line with the abscissa, usually referred to as the virtual origin. The straight-line behaviour does not hold in the developing region (approximately for $x/d < 5$) [11, p. 100].

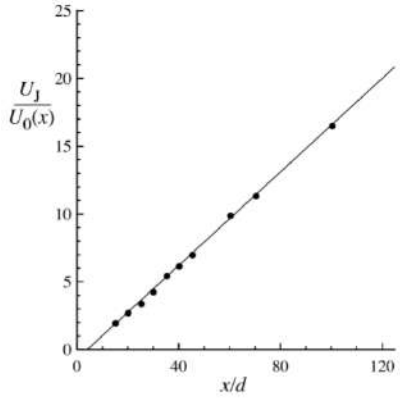

Figure 1.3: Variation of the scaled centerline velocity $U_J/U_0(x)$ as a function of the nondimensional distance x/d. Reproduced from [11].

Consistent with the decay of the centerline velocity, it is observed that the jet spreads linearly in the axial direction. Because the flow is turbulent, we do not consider instantaneous flow values; instead, we work with time-averaged quantities, denoted by the

averaging brackets $\langle\cdot\rangle$. In cylindrical coordinates, the mean axial velocity field is written as $\langle U(x, r, \theta)\rangle$. The jet half-width $r_{1/2}(x)$ is then defined as the radial location where $\langle U(x, r_{1/2}, 0)\rangle = \frac{1}{2}U_0$, and experimental data show that it follows the empirical law:

$$r_{1/2}(x) = S(x - x_0) \tag{1.5}$$

where S is the constant spreading rate; this linear spreading is a direct consequence of the overall continuity of the flow and decay of the centerline velocity [11, p. 100].

Substituting the centerline decay (Eq. 1.4) and the linear spreading (Eq. 1.5) into the time-averaged BLE and introducing a turbulent viscosity ν_T, a closed-form solution for the mean radial velocity profile can be derived [11, p. 119]. In terms of the similarity variable $\eta = r/(x - x_0)$, the solution is:

$$f(\eta) = \frac{\langle U\rangle}{U_0} = \frac{1}{(1 + a\eta^2)^2} \tag{1.6}$$

where the coefficient a is related to the spreading rate S by:

$$a = \frac{\sqrt{2} - 1}{S^2} \tag{1.7}$$

A key consequence of this result is self-similarity: although the profile evolves along the downstream distance, its shape remains unchanged when scaled by U_0 and $r_{1/2}$, so that profiles at different axial locations all collapse onto a single curve (see Figure 1.4).

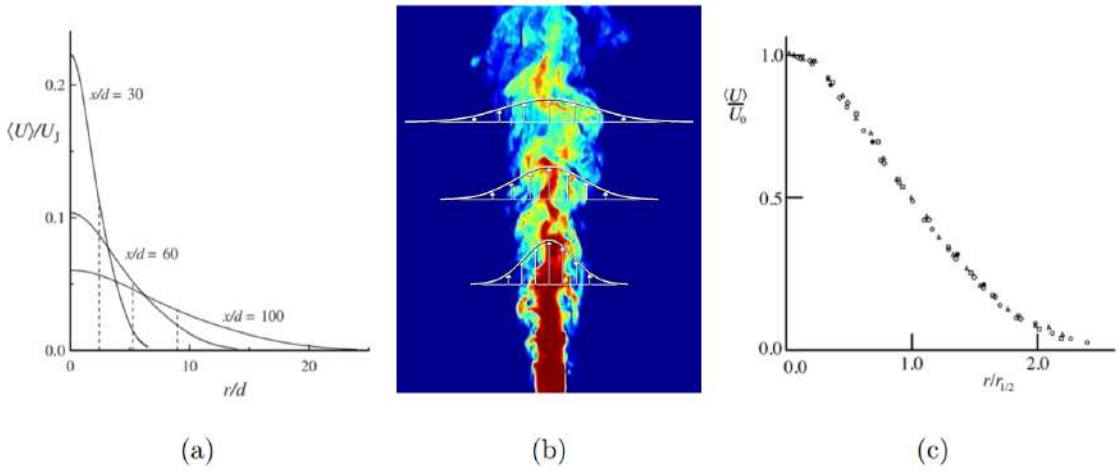

(a) (b) (c)

Figure 1.4: (a) Radial profiles of mean axial velocity, the dashed lines indicate the half-width. Reproduced from [11]. (b) Qualitative representation of the radial profiles of the mean axial velocity $\langle U\rangle$ of a turbulent jet flow. Adopted from [15]. (c) Scaled versions of the radial profiles of mean axial velocity; different symbols represent different axial distances. Reproduced from [11]

1.3 Particle Tracking Velocimetry

To track snow particles through the atmosphere or the motion of tracer particles in a jet, this thesis employs Particle Tracking Velocimetry (PTV) to reconstruct the motion of individual tracer particles from time-resolved images. PTV, also referred to as Lagrangian Particle Tracking (LPT), is a measurement technique used to track the three-dimensional motion of individual tracer particles in a fluid. By reconstructing particle trajectories over time, PTV provides direct access to Lagrangian statistics of the particle motion, while also enabling the reconstruction of Eulerian flow fields from the same data (section 1.5).

PTV is divided into three main tasks: localisation of particle centres in the two-dimensional camera frames, three-dimensional reconstruction of particle positions, and creation of tracks using sequences of images.

Localisation In the first step, the goal is to identify the centre of each particle image from the camera-recorded intensity peaks, with sub-pixel accuracy. The detection algorithm must also be robust to image noise and capable of handling partially overlapping particle images. A simple way to do that is to use a weighted average of the pixel intensities, but Gaussian fitting of the intensity peaks is usually preferred, as it provides more robust and accurate sub-pixel localisation [16].

Reconstruction In the second step, the particle centres detected in the individual cameras are combined to find the three-dimensional location of each particle. This can be done by triangulation, a geometric process that finds the three-dimensional point that best fits the viewing rays defined by the previously calibrated cameras. More advanced iterative reconstruction methods will be discussed in section 1.4 [17].

Tracking In the third step, the reconstructed particle positions at successive time steps are linked to form continuous trajectories. An intuitive way to perform this task is to associate each particle with its nearest neighbour in two consecutive frames. However, this approach becomes problematic at high particle densities when the distance between neighbouring particles can be smaller than the displacement occurring between two frames. To reduce these ambiguities, multi-frame schemes such as the 4-frame best-estimate method use information from several consecutive frames to predict particle motion and enforce temporal consistency [16].

1.4 Shake the Box

Shake The Box (STB) is a high-density LPT method that was originally developed to improve the performance of Tomographic PIV [18] and classical PTV [19] at high seeding

densities. Instead of reconstructing a full volumetric intensity field at each time step, STB directly iterates on a sparse set of particle positions and trajectories, strongly exploiting temporal coherence between consecutive images [20].

The STB algorithm can be conceptually divided into three phases: initialisation, convergence, and the converged state. In the following, these stages are briefly outlined. Figure 1.5 gives an overview of the workflow.

Initialisation phase In the first images of a dataset, no *a priori* track information is available to guide the prediction. The aim of the initialisation phase is therefore to identify as many physically consistent tracks as possible over a few time steps while discarding ghost particles and spurious detections. In practice, this stage is typically applied to the first four frames of each sequence [20][21].

Convergence phase The tracks identified during the initialisation phase are then used to predict particle positions at subsequent time steps. Based on these predictions, additional tracks can be established and ghost tracks are progressively removed, until a stable set of physically consistent trajectories is obtained. Once the particle positions have been predicted, their locations are refined by the so-called "shaking" procedure. Each predicted particle is slightly displaced in three dimensions so that its synthetic image, obtained from the current optical model, best matches the particle image measured by the cameras. By subtracting the synthetic particle images from the original recordings, residual images are obtained that highlight remaining intensity peaks not yet selected as particles. These residual peaks are then used to introduce new particles and to initiate additional tracks, further improving the completeness and consistency of the reconstruction [20][21].

Converged phase In the converged phase, the number of tracked particles no longer changes significantly, and a quasi-steady state is reached. The processing workflow remains essentially the same as in the convergence phase, but most trajectories are now extended over many time steps and typically terminate only when the corresponding particle leaves the measurement volume. Conversely, new tracks are predominantly created for particles that enter the illuminated volume at later times [20][21].

The accuracy of the STB algorithm strongly depends on the quality of the three-dimensional reconstruction of the particle positions. To achieve such accuracy at high particle image densities, STB is typically combined with the IPR technique introduced by Wieneke [22]. In IPR, an initial triangulation is followed by iterative steps in which the current 3D particle ensemble is reprojected onto the image planes to generate synthetic intensities. These synthetic images are compared to the measured ones, and the resulting

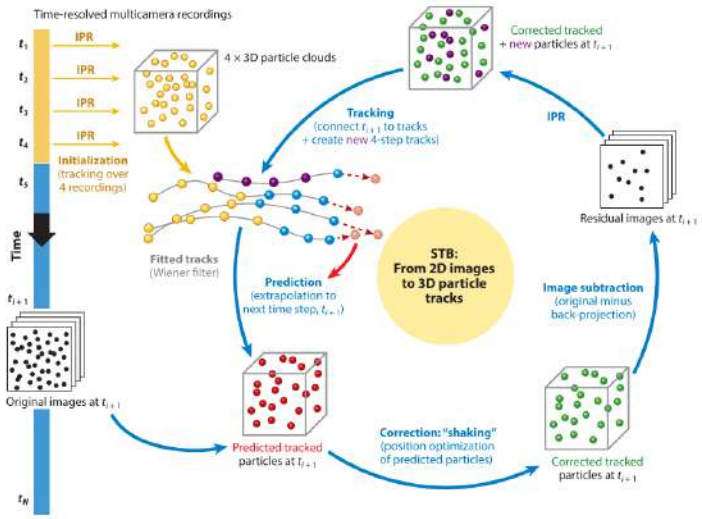

Figure 1.5: Visual interpretation of STB. The initialisation phase (with Iterative Particle Reconstruction (IPR)-based particle detection) is highlighted in yellow, while the convergence and converged phases are shown in blue. Reproduced from [21].

residuals are used to adjust particle positions and detect previously unresolved intensity peaks, enabling sub-pixel accuracy even for high-density recordings.

Thanks to recent developments in the STB algorithm and advanced IPR techniques [23], STB can now converge for particle image densities of up to about 0.2 Particles-per-pixel (ppp) in synthetic datasets with negligible noise, and up to approximately 0.15 ppp in experimental recordings under well-controlled conditions. These values mark a substantial improvement over conventional PTV, which is typically limited to particle image densities of about 0.005 ppp, and are crucial for exploiting the high-seeding conditions used in the present work [24].

1.5 FlowFit

In the analysis of turbulent flows, it is often essential to also obtain an Eulerian description of the motion to access canonical quantities defined on a fixed spatial grid, such as vorticity or the Q-criterion, and to relate the measurements to well-established flow structures and stability concepts [20].

Traditionally, this Eulerian representation is obtained by interpolating the Lagrangian quantities associated with the particle tracks onto a regular grid, for instance, by averaging the velocities of all particles contained in each cell or by using distance-weighted kernels to distribute their contribution in space. However, this purely geometric technique is very sensitive to local variations in particle density and measurement noise, which can strongly affect the quality of the result.

These limitations can be overcome by the physics-based data assimilation method FlowFit3 (FF3) [25]. FF3 projects the discrete and discontinuous Lagrangian measurements onto three-dimensional B-splines, which act as a low-pass filter to reduce noise and can be evaluated efficiently at arbitrary positions within the reconstructed volume. Moreover, FF3 is designed to enforce physical consistency by using the Navier–Stokes equations as a constraint and by imposing incompressibility, minimising the associated gradient-based error functional. All of this enables a faster and more robust reconstruction of the flow in an Eulerian framework, significantly improving the quality and reliability of the resulting fields [26].

1.6 Goal and Outline

The main goal of this thesis is to investigate Lagrangian particle tracking in two configurations, namely natural snowfall in the atmospheric boundary layer and a turbulent round jet in the laboratory. In particular, the work aims to compare these two measurement approaches and to analyse their differences and interconnections, with the perspective of advancing large-volume LPT by combining field-scale observations with the accurate laboratory methods and algorithms provided by STB.

Chapter 2 describes the experimental setups used in the field campaign in Davos, Switzerland, and in the laboratory experiment at German Aerospace Center (DLR)[1]. **Chapter 3** presents the data-processing pipelines for both configurations, from camera calibration to track reconstruction. **Chapter 4** reports the main results, first for the snowfall measurements and then for the water-jet experiments, considering both Lagrangian and Eulerian perspectives. **Chapter 5** discusses these results, comparing the two approaches and relating them to existing literature. **Chapter 6** summarises the main conclusions of this work and provides recommendations for future developments.

[1]Deutsches Zentrum für Luft- und Raumfahrt

Chapter 2

In the Field and in the Lab

To realise measurements of natural snowfall and a water jet, two experimental setups are described in this work. The first is an outdoor system for tracking natural snowfall in the atmospheric boundary layer, while the second is a laboratory configuration for studying a turbulent round jet under well-controlled conditions.

2.1 Outdoor Snow Tracking Experiment

To track natural snowfalls, the measurements are conducted in Davos, Switzerland, at a field site near the ski slopes on the Weissfluhjoch, at an elevation of 2503 m above sea level during the 2025–2026 winter campaign. The site is accessed by taking a funicular railway to the top station, followed by a descent of a few hundred metres either on snowshoes or skis, which makes transporting all the equipment required for the measurements particularly demanding; see Figure 2.1.



Figure 2.1: Overview of the experimental site in Davos, Switzerland. The tents are deployed in a 30-meter diameter arc around the light panels.

2.1.1 Experimental Layout

Once at the site, all equipment was first secured inside a closed and partially heated hut, after which the deployment of the setup began. Because of the harsh weather conditions and the limited time available, this phase had to be carefully planned in advance to ensure that everything was prepared correctly and that the system was ready for measurements at the predefined acquisition times.

Figure 2.1 provides an overview of the measurement site and its surroundings. The seven tents can be seen next to the hut, deployed along a 30 m arc around the 10 $\times$ 10 $\times$ 10 m^3 measurement volume, which starts approximately 2 m above the ground. The tents are arranged such that four of them are blue (every other tent), the central one is orange, and the remaining two are red.

Each tent hosts a specific part of the imaging system and the associated hardware, protecting against low temperatures, wind, and snow; each blue tent hosts four cameras mounted on a 2x2 array that is positioned on a tripod, imaging a small triangular opening. The red tents contain the computers and batteries needed to operate the camera system during the measurements. The orange tent was used as a central control station for coordinating the campaign and for data acquisition (network tent). All tents had to be firmly anchored to the ground and reinforced with snow walls on their sides, as the strong winds could blow away tents that were not properly secured.

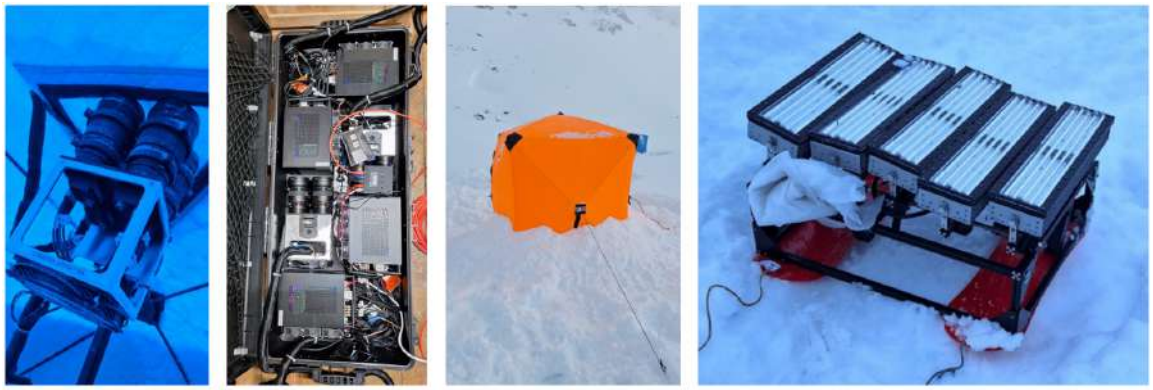

Figure 2.2: Main components of the volumetric imaging system. (a) Camera setup: 16 cameras were used in total, divided into 4 arrays of 4 cameras, located in the blue tents. (b) Control computer used for synchronising the cameras and recording the volumetric snowfall data. (c) Tent setup: 7 tents in total were deployed, protecting the cameras and computers during data acquisition. (d) LED light panels providing the illumination necessary to detect snowflakes.

At the centre of the tent arc, directly below the measurement volume, five light panels were deployed in a straight line. They ensured adequate illumination of the snowflakes

traversing the volume, enabling reliable detection by the cameras. Each of the light panel units was composed of 5 white LED modules, mounted on small snowboards and pointing upwards. The panels, cameras, and remaining hardware were powered by batteries that supplied the required energy, and all components were interconnected with long cables for data acquisition. This led to a challenging coordination problem during deployment and made moving around the site difficult, as damaging a single cable could compromise the entire measurement. A close-up view of the camera arrangement, tents, and light panels described above is shown in Figure 2.2.

2.1.2 Drone-based Camera Calibration

After the deployment of the site, a calibration flight was performed to determine the projective geometry of the imaging system. This was done using a DJI Matrice 300 drone, on which four light bulbs were mounted in a tetrahedral configuration, as shown in Figure 2.3b. The drone followed a pre-programmed flight path, moving horizontally and vertically through multiple planes of the measurement volume to sample it as uniformly and completely as possible. The recorded images of the moving lights were then used to estimate the projection matrices of all cameras and to refine their relative positions and orientations, see [2]. The details of the calibration procedure and its accuracy are further discussed in Section 3.1.1.

2.1.3 Snowfall Measurement Campaign

During the winters of 2024/2025 and 2025/2026, several snowfall imaging campaigns were carried out at the Tschuggen and Weissfluhjoch field sites, using both volumetric and planar configurations of the large-scale setup. In the volumetric configuration, cameras were arranged around the illuminated volume to enable three-dimensional reconstruction, whereas in the planar configuration, they were aligned approximately along a line and imaged in the same direction, such that a two-dimensional measurement plane could be reconstructed instead of a volume.

Table 2.1 summarises all snowfall imaging campaigns, listing the date, location, number of datasets, measurement configuration (planar or volumetric), and the type of calibration performed. The volumetric calibrations were carried out either using the DJI Matrice 300 drone introduced in Section 2.1.2, or with a DJI Mavic 3E drone equipped with a single light source.

The description and photographs in Chapter 2.1 refer to the field deployments at Weissfluhjoch on 09/01/2026 and 10/01/2026. In contrast, the snowfall datasets processed in this work and analysed in Chapters 3, 4, and 5 were acquired earlier at Tschuggen during the campaigns of 12/03/2025 and 13/03/2025.

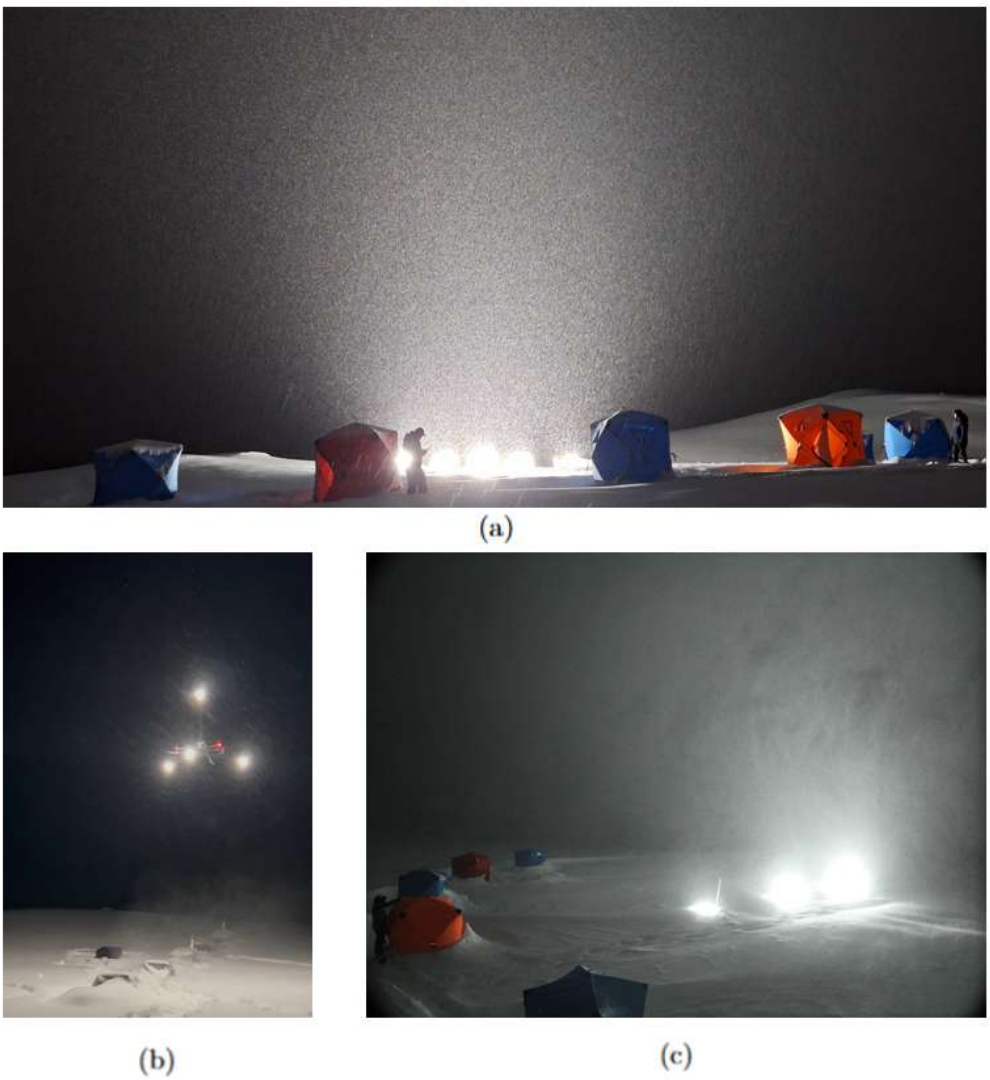

Figure 2.3: (a) Large scale snow tracking experiment in low wind condition. (b) Calibration flight. The four light bulbs mounted on the drone follow a predetermined path through the measurement volume and are recorded by the cameras for calibration purposes. (c) Overview of the illuminated measurement volume and tent arrangement during high wind conditions.

2.2 Laboratory Water Jet Tracking Experiment

In contrast to the outdoor snow-tracking experiment described in Section 2.1, this section presents a laboratory water-jet experiment carried out at DLR. This setup aims to study a turbulent round jet under well-controlled conditions, using a volumetric imaging system to visualise the flow at different Reynolds numbers.

2.2.1 Experimental Layout

A schematic overview of the experimental setup is shown in Figure 2.4a. The water jet is generated at the bottom of a hexadecagonal-sided tank (with a diameter of $73cm$ and a

Table 2.1: Overview of the snowfall imaging datasets measured during the 2024/2025 and 2025/2026 campaigns. For each date, the type and number of datasets, as well as which drone has been used for the calibration, are reported.

Date	Location	Number of Datasets	Measurement Type	Calibration
14/12/2024	Tschuggen	2	Planar	Mavic
27/01/2025	Tschuggen	1	Planar	Mavic
12/03/2025	Tschuggen	4	Volumetric	Mavic
13/03/2025	Tschuggen	5	Volumetric	Mavic + Matrice
14/04/2025	Weissfluhjoch	1	Volumetric	Matrice
09/01/2026	Weissfluhjoch	3	Volumetric	Mavic + Matrice
10/01/2026	Weissfluhjoch	5	Volumetric	Mavic + Matrice
12/02/2026	Weissfluhjoch	5	Volumetric	Mavic + Matrice
14/03/2026	Tschuggen	5	Planar	Mavic

height of 68cm), where water is injected through a round 10 mm nozzle and flows vertically upwards into the surrounding fluid. The water in the tank is pumped into a reservoir located a few metres higher, from which it flows back down to the nozzle, forming a closed loop. Along this loop, the flow passes through a flow metre that allows the jet exit velocity to be monitored and controlled.

The water is seeded with tracer particles that have a diameter of approximately 50μm, which are detected by four cameras arranged around the tank as they are illuminated by a light beam entering the measurement volume from above. Everything is mounted on rigid aluminium profiles and posts, which provide a stable and adjustable framework for the precise alignment of the volumetric imaging system.

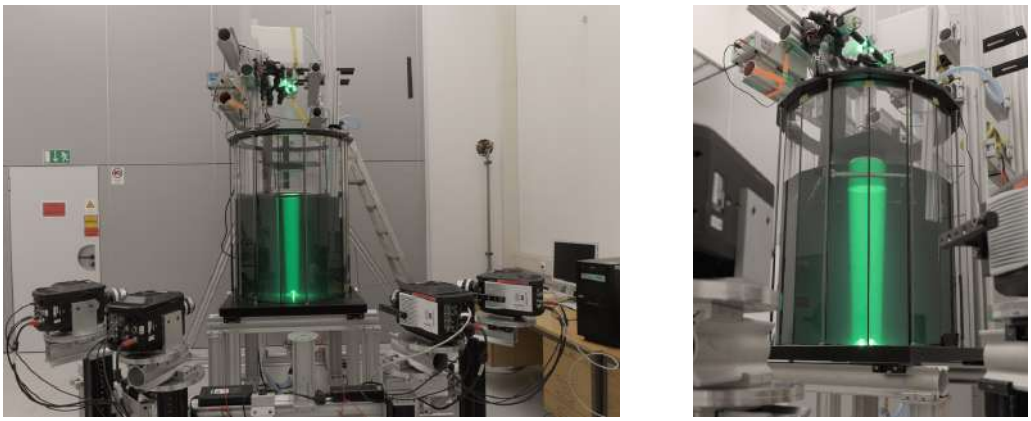

(a) (b)

Figure 2.4: (a) General setup of the water jet experiment. Four cameras are positioned around the hexadecagonal-sided tank, pointing at the light beam to capture illuminated tracer particles emerging from the jet nozzle at the bottom. (b) Close-up view of the illuminated measurement volume inside the tank, showing the light beam generated by the optical system.

2.2.2 Illumination and Optical System

The light beam used to illuminate the tracer particles is generated by a dedicated LED-based optical system, shown in Figure 2.5c. To obtain reliable measurements, a strong illumination is required so that the tracer particles are clearly visible in the images. To increase the available light power in this experiment, two in-house assembled LED units are used, both directed towards a prism that redirects the individual beams into the desired direction and superimposes them into a single, more intense light beam.

At this point, the light beam is still relatively broad, whereas the desired diameter of the measurement volume is approximately $8cm$. To obtain this diameter, a converging lens (see Figure 2.5a) is placed about $50cm$ below the prism, which not only shapes the beam to the required size but also increases its intensity.

Lastly, the optical setup includes two additional components that improve the quality of the illumination. First, a glass panel lying on the water surface is used to maintain a flat interface and reduce refractions caused by jet-generated surface waves (see Figure 2.5a). Second, an aperture mounted around the lens blocks the portion of the light beam that lies outside the optical path, restricting the illumination to the measurement volume and enhancing the contrast of the tracer particles (see Figure 2.5b).

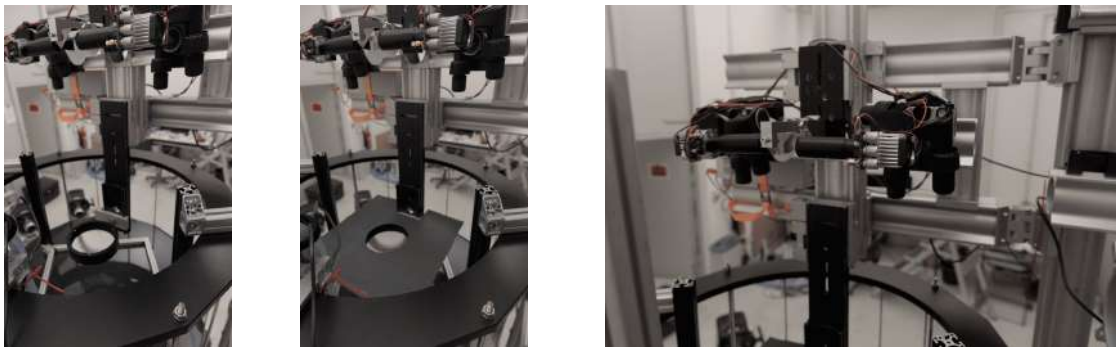

(a) (b) (c)

Figure 2.5: Optical components. (a) Converging lens positioned above the water surface, with a glass panel underneath to eliminate refraction caused by surface waves generated by the flow. (b) Aperture (passepartout) mounted on top of the lens to block stray light outside the optical path. (c) Optical setup for generating the light beam. Light from two LEDs is combined using a prism to create a uniform illumination volume for particle tracking.

2.2.3 Camera Arrangement and Calibration

To measure the light reflected by the tracer particles, four cameras are arranged around the tank, each viewing the illuminated measurement volume from a different angle, as shown in Figure 2.6a. When comparing this figure with Figures 2.3c, it is noteworthy that

the water jet cameras are in a similar configuration to the snow tracking experiments.

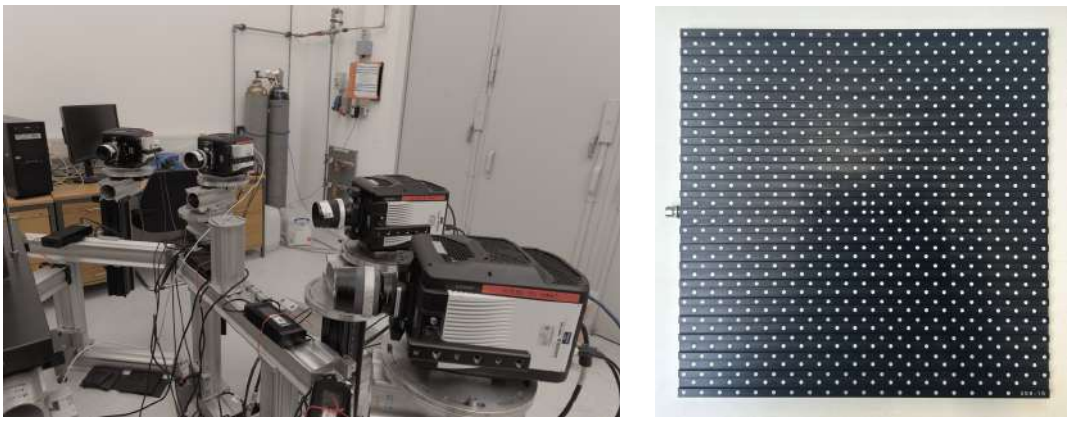

(a) (b)

Figure 2.6: (a) Close-up view of the four-camera array positioned around the measurement volume. (b) Freely moving calibration target with markers distributed on two planes to provide a shallow depth.

The first step after mounting the setup is to perform the calibration. In this experiment, a freely moving target is used, conceptually similar to the approach of [27], in contrast to the predefined drone path described in Section 2.1. The target consists of a rigid plate carrying a well-defined pattern of circular markers with known spacing (see Figure 2.6b). It is manually moved through the measurement volume at different positions and orientations so that the pattern is visible to all cameras at multiple locations, thereby covering the entire Field Of View (FOV). During this procedure, image sequences of the target are recorded simultaneously by the four cameras.

2.2.4 Water Jet Measurement Campaign

Once the setup is aligned, the operating conditions are fixed, and the calibration is done, measurements are performed for several flow regimes. Multiple datasets are acquired at three constant Reynolds numbers, corresponding to an early transitional state at $Re = 2000$, a late transitional state at $Re = 4000$, and a fully turbulent state at $Re = 11000$. In addition, one dataset is recorded with a linearly increasing Reynolds number, starting from $Re = 2000$ and reaching $Re = 11000$ over a period of 27 seconds. The two transitional cases are recorded at a camera frame rate of $1000Hz$, whereas the turbulent and ramped-Re datasets are acquired at $2000Hz$.

Figure 2.4b shows the experimental setup during data acquisition, with the measurement volume recorded by the cameras extending from the bottom of the water tank to approximately 30 cm above the nozzle.

Chapter 3

Data Processing and Calibration

This chapter describes the processing and calibration procedures applied to both the outdoor snow-tracking campaign and the laboratory water-jet experiment. First, we present the drone calibration and the reconstruction of snowflake trajectories in the atmospheric boundary layer; we then introduce the corresponding procedures for the tracer particles in the turbulent water jet. Together, these sections provide the methodological foundation for the Lagrangian and Eulerian flow analyses discussed in Chapter 4.

3.1 Snow Tracking

The analysis of the snow-tracking measurements relies on a dedicated calibration and processing chain that converts the raw multi-camera images into Lagrangian particle trajectories suitable for quantitative flow analysis. This section presents the main steps of this pipeline and outlines how they are applied to the present dataset acquired in March 2025; see Table 2.1.

3.1.1 Drone Calibration Framework

The analysis of the snow-tracking measurements starts with a dedicated camera-array calibration framework based on a drone-borne light marker system. In this section, we briefly summarise the main ingredients of this framework, which establishes the mapping between the three-dimensional calibration markers on the drone and their two-dimensional projections in all cameras. The corresponding mathematical formulation provides the basis for the subsequent reconstruction of snowflake trajectories introduced in Section 4.2.

To calibrate the camera array, we define a mapping between a point in object space $\boldsymbol{X} = [X\ Y\ Z]^T$ and its projected image coordinate $\boldsymbol{x} = [x\ y]^T$ in the camera reference frame. Following the framework currently developed at Institute of Fluid Dynamics (IFD), this mapping is written in augmented form $\tilde{\boldsymbol{x}} = [x\ y\ k]^T$, where the additional parameter

k is aligned with the camera depth of field [17, p. 2].

The complete projection model is then expressed by the following compact equation:

$$d(K^c\tilde{\mathbf{x}}) = \underbrace{R_k(\theta^c)\begin{bmatrix} I^\star & \begin{bmatrix}\mathbf{t}^c\\ 0\end{bmatrix}\end{bmatrix}\begin{bmatrix} R^a & \mathbf{T}^a \\ \mathbf{0}^T & 1\end{bmatrix}}_{P^{a,c}}\begin{bmatrix}\mathbf{X}\\ 1\end{bmatrix} \text{ with } K^\star = \begin{bmatrix}\alpha_x & s & \alpha_x p_x \\ 0 & \alpha_y & \alpha_y p_y \\ 0 & 0 & f\end{bmatrix} \text{ and } I^\star = \begin{bmatrix}0 & 1 & 0\\ 1 & 0 & 0\\ 0 & 0 & 1\end{bmatrix}. \tag{3.1}$$

Here $K^\star$ defines the (inverse) intrinsic camera calibration matrix with pixel pitch $\alpha_{x,y}$, focal length f, image centre $p_{x,y}$, and skew s. Furthermore, R^a and $\mathbf{T}^a$ are the array rotation matrix and translation vector, where the rotation $R_k(\theta^c)$ with angle θ^c and translation vector $\mathbf{t}^c$ encode the relative spacing and rotation for the individual cameras in the array. Together, these define $P^{c,a}$, the extrinsic projection matrices, with $I^\star$ a unitary identity that flips the axes.

The projection model in Equation 3.1 is based on an ideal pinhole geometry and, therefore, does not account for lens distortion. In the present work, optical distortion is represented by a scalar distortion function $d(\bullet)$ based on a first-order refractive interface correction modelled from Snell's law to the lens optics [28] with index parameter b and applied to the camera reference:

$$d\left(\tilde{\mathbf{x}}; b\right) = \frac{1}{\sqrt{b(x^2+y^2)+k^2}}\begin{bmatrix}x\\ y\\ k^*\end{bmatrix}, \tag{3.2}$$

where $k^* = \sqrt{b(x^2+y^2)+k^2}$ enforces the augmentation unitary.

In the present calibration, the three-dimensional points $\boldsymbol{X}$ represent the world-space positions of the light markers mounted on the drone, while the corresponding two-dimensional points $\boldsymbol{x}$ denote their image coordinates in each camera. By recording the drone at multiple positions and orientations along its trajectory (see Figure 3.1a), we obtain a large set of 3D–2D correspondences over many frames, which are then used to estimate the projection matrices by solving the following minimisation problem:

$$\min_{\{k_i^{c,n} \;\cdots\; \mathbf{T}_d^n\}} \sum_{i,n,c,a} \left\| d\left(K^c\tilde{\mathbf{x}}_i; \tilde{n}\right) k_i^{c,n} - P^{a,c}\begin{bmatrix} R_d^n & \mathbf{T}_d^n \\ \mathbf{0}^T & 1\end{bmatrix}\begin{bmatrix}\mathbf{X}_i\\ 1\end{bmatrix}\right\|_2^2, \tag{3.3}$$

for which we introduce the rotation matrix R_d^n and the translation vector $\mathbf{T}_d^n$ that define the drone's position and orientation. Summing over all cameras and calibration frames, this defines a nonlinear global residual error minimisation that optimises the tracking of the drone position with the best possible accuracy and can be numerically solved with

methods described in [29, Chap. 9].

An alternative way to relate three-dimensional object-space coordinates to two-dimensional image coordinates, while explicitly accounting for distortion, is provided by the calibration mapping proposed by Soloff et al. [30]. In this framework, the image-plane coordinates are expressed as a truncated nonlinear polynomial basis function expansion of the object-space coordinates.

$$\boldsymbol{x} = \boldsymbol{F}(\boldsymbol{X}) \tag{3.4}$$

where $\boldsymbol{F}$ is a nonlinear distortion-compensated mapping identified from calibration data.

In practice, the mapping function $\boldsymbol{F}$ is not known but can be represented by a suitable polynomial function, flexible enough to capture distortion patterns while still allowing for an efficient least-squares identification of its coefficients from calibration data.

$$\hat{\boldsymbol{F}}(\boldsymbol{X}) = \sum_k \boldsymbol{a}_k \phi_k(\boldsymbol{X}) \tag{3.5}$$

where $\phi_k \in \{X, Y, Z, X^2, Y^2, Z^2, XY, ...\}$ denotes the basis functions and $\boldsymbol{a}_k = \begin{bmatrix} a_1 & a_2 \end{bmatrix}_k^T$ the corresponding coefficients. From a mathematical point of view, this corresponds to retaining only a finite number of terms in a Taylor expansion of the true distortion mapping around the calibration region, so that higher-order terms provide only small corrections and can be neglected.

In their original work, Soloff et al. [30] demonstrate that a cubic polynomial already yields very small residual errors for stereoscopic PIV configurations, suggesting that relatively low-order polynomials are sufficient in many practical cases. The specific polynomial representation adopted in this study, as well as its impact on the calibration accuracy and

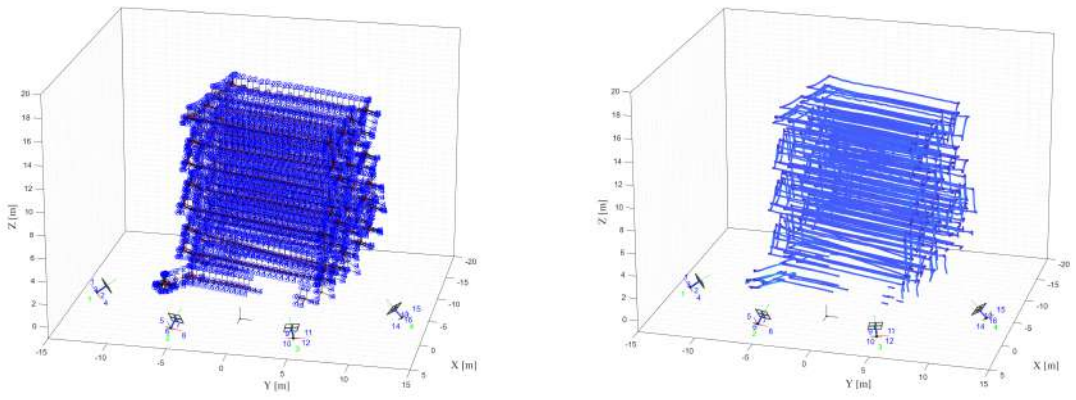

(a) (b)

Figure 3.1: (a) Drone trajectory used for the calibration flight, reconstructed in MATLAB from the three-dimensional coordinates over the measurement volume. (b) Drone velocity along the path.

subsequent flow measurements, will be discussed together with the snow-tracking results in Chapter 4.

3.1.2 Snow Tracking Reconstruction

Once the cameras are calibrated, the snowflakes recorded by the 16 cameras must be reconstructed in the three-dimensional world frame and then tracked over successive frames. This allows us to recover their Lagrangian trajectories as accurately as possible, maximising the number of physically consistent tracks while avoiding spurious reconstructions and ghost particles.

The code used to perform this task was developed in-house at the IFD and is based on the tracking framework proposed by Muller in [31] for schooling fish. It is a robust Lagrangian particle tracking framework tailored to large multi-camera volumes. The dedicated image pre-processing, background suppression, and accurate multi-view triangulation enable the reconstruction of physically consistent trajectories, even under varying illumination and partial occlusions.

In the present application, this framework is directly applied to illuminated snowflakes in the atmospheric boundary layer, with minimal adjustments to the underlying tracking logic originally developed for schooling fish. For further reading, see the previous IFD-thesis [2] [32] for the initial developments of snow imaging and tracking.

3.2 Turbulent Water Jet

The analysis of the turbulent water-jet experiment is based on a dedicated processing workflow that converts the raw volumetric images into three-dimensional particle trajectories and Eulerian flow fields. This section outlines the key stages of this workflow, including camera calibration, volume self-calibration, and STB processing.

3.2.1 Calibration using Target

For successful LPT with STB, the first step after image acquisition is to calibrate the camera system. The calibration provides the mapping between the two-dimensional particle images recorded by each camera and the corresponding three-dimensional particle positions in the measurement volume.

In this experiment, the camera–volume mapping is obtained from images of a freely moving calibration target, as described in Section 2.2, so that the entire FOV is sampled and each camera observes the target from multiple viewing angles. Each camera detects the marker positions and fits the corresponding calibration grid to the recorded image (see Figure 3.2a). The measurement volume is then represented by 13 calibration planes along

the depth, and for each plane, a third-order polynomial mapping is used to transform world coordinates (x' and y') on that plane into the corresponding raw image coordinates (x and y); an example of this polynomial mapping is shown in Figure 3.2b.

The overall procedure is conceptually equivalent to the large-scale multi-camera calibration method of Muller et al. [27], in which a planar chequerboard with known internal geometry is moved freely to sample the full field of view and recover both the camera locations and their effective focal lengths.

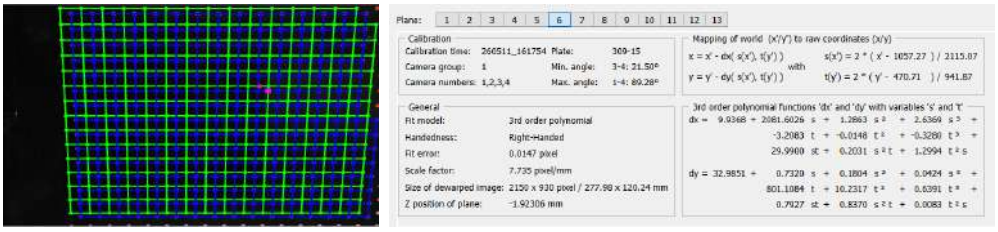

(a) (b)

Figure 3.2: Camera calibration in DaVis using a freely moved target. (a) Target image recorded by one camera (Camera 4), with the fitted calibration grid overlaid on the detected marker points. (b) Corresponding calibration mapping parameters for plane 6 (centre of the measurement volume).

This first calibration achieves a very high accuracy, with a mean error on the order of hundredths of a pixel (0.0147 for plane 6, see Figure 3.2b), indicating that the third-order polynomial mapping reproduces the target markers almost perfectly. Nevertheless, this accuracy is expected to deteriorate for the actual flow measurements, as small misalignments between the calibrated projection model and the true optical paths can deviate and accumulate over time due to refractive-index inhomogeneities, mechanical shifts of the setup, or limitations of the polynomial distortion model, which is fitted only to a finite set of sparse target markers.

3.2.2 Volume Self Calibration

Usually, calibrating the cameras using only a target is not sufficient. To achieve sub-pixel accuracy, a Volume Self Calibration (VSC) is applied using several images of the tracer particles [33].

During this process, the measurement volume is divided into subvolumes; in the present experiment, the depth (z) direction is split into 13 layers, while the other two directions are divided into 5 subparts each. For each subvolume, the program selects a large number of particle image intensity peaks (about 90'000 in total), reconstructs their 3D positions using the current calibration, and reprojects them onto the camera images. The distance between the reprojected location and the corresponding measured peak on each camera defines a local disparity, which is accumulated within each subvolume and visualised in a disparity map (see Figure 3.3a). For a well-calibrated system, the disparity map in each

subvolume collapses into a small, approximately circular spot centred at zero displacement, indicating that the reprojected particle positions coincide with the measured image peaks.

For each subvolume, a disparity vector is obtained by averaging the reprojection errors of all particles contained in that region. The collection of these vectors over the entire measurement volume defines a (quasi-)continuous disparity field (see Figure 3.3b), which is then used to update the mapping parameters that transform the reconstructed 3D particle coordinates into the corresponding image coordinates.

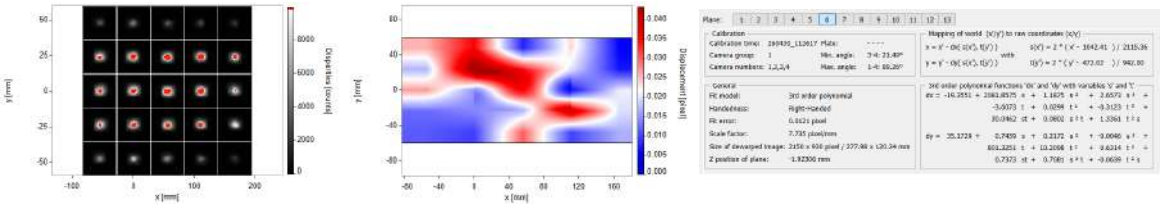

(a) (b) (c)

Figure 3.3: VSC in DaVis. (a) Disparity map obtained from particle images on plane 6. (b) Corresponding disparity field used to update the camera-volume mapping. (c) Updated mapping parameters after VSC.

This calibration step compensates for the small residual misalignments left by the first, target-based calibration, reducing the mean disparity from values of the order of several tenths of a pixel to values of the order of a few hundredths of a pixel and thus providing favourable conditions for the subsequent STB processing. In principle, the procedure can be applied iteratively, but in practice, one or two VSC iterations are usually sufficient to reach the desired sub-pixel accuracy.

3.2.3 Optical Transfer Function

In addition to the geometric calibration, which provides an accurate 3D position for each particle, an Optical Transfer Function (OTF) is reconstructed from low particle–density images, as shown in Figure 3.4a. This step is crucial because STB relies on comparing reprojected synthetic particle images with the recorded ones. Therefore, not only must the particle position be represented accurately, but its image shape must also be accurate. The OTF characterises how an ideal point-like tracer appears on the sensor, its shape being embedded in the recorded images, and thus defines the particle image kernel used during the STB shaking procedure. [33]

The main principle is that a spatially homogeneous mapping between voxels and image intensities cannot be assumed. Optical effects such as blurring caused by the finite depth of field and astigmatism introduced by the imaging optics vary across the field of view and therefore cannot be neglected. The measurement volume is divided into subvolumes in the same way as for the VSC described in the previous section. Within each subvolume,

the particle image shapes are characterised via the OTF and then averaged to obtain a representative particle image for that region. This averaged shape is used as the particle image model when the 3D particle positions are reprojected into the image planes, for example, when computing the residual images during the STB processing. [34]

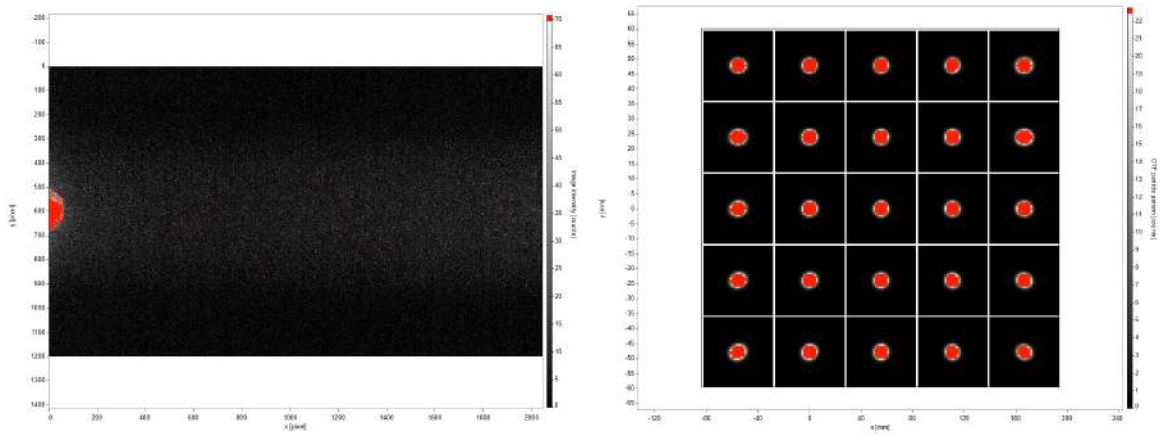

(a) (b)

Figure 3.4: OTF reconstruction in DaVis using low particle density images. (a) Low-density particle image recorded by one camera. (b) Corresponding OTF particle pattern for plane 6.

Figure 3.4b shows the particle patterns for all subvolumes in one calibration plane along the depth. It can be seen that the particles are generally round, as expected, but some distortion appears towards the edges of the field of view, where the particle images become slightly elongated and elliptical.

Thanks to this calibration, the synthetic particles used by STB closely match the recorded ones not only in position but also in shape, which improves the sensitivity of the residual images to small changes in particle location and thus supports a more robust convergence of the STB reconstruction.

3.2.4 Commercial STB DaVis

Once the calibration has been completed, the commercial implementation of STB in DaVis (from LaVision) has been used to reconstruct three-dimensional particle tracks from the water jet recordings.

Before processing, some preprocessing is usually applied to the raw images to increase the contrast and help the algorithm detect particles more reliably. However, such operations can also remove information from the particle images and introduce biases into the measurements, leading to results that are less representative of the true flow. For this reason, only minimal preprocessing is used in the present work, limited to background subtraction, which does not alter the physical behaviour of the flow; this approach is

only feasible thanks to the high signal quality and geometric precision achieved in the laboratory experiments.

Figure 3.5a shows a representative raw image of the experiment: the jet nozzle is visible on the left, and the measurement volume is densely filled with tracer particles. In Figure 3.5b, the minimum intensity recorded at each pixel over 2000 frames is displayed; the jet appears very bright because the nozzle remains essentially stationary and continuously reflects light. A band of low-intensity pixels can also be observed near the end of the measurement volume and around the water jet, due to light scattering and non-uniform illumination.

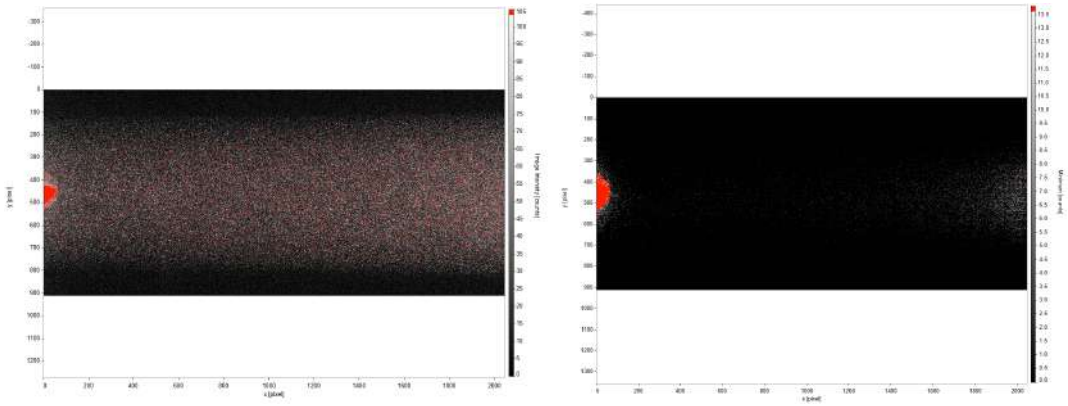

(a) (b)

Figure 3.5: Image preprocessing in DaVis. (a) Raw particle image before minimum subtraction. (b) Minimum intensity image subtracted from the recorded sequence. Note that the colour scales differ between the two panels.

The minimum-intensity image is then multiplied by a factor of 1.3 to make the background subtraction slightly more aggressive and subsequently subtracted from each raw frame. The resulting images, with most of the stationary reflections and low-frequency background removed, are then used as input for the STB processing.

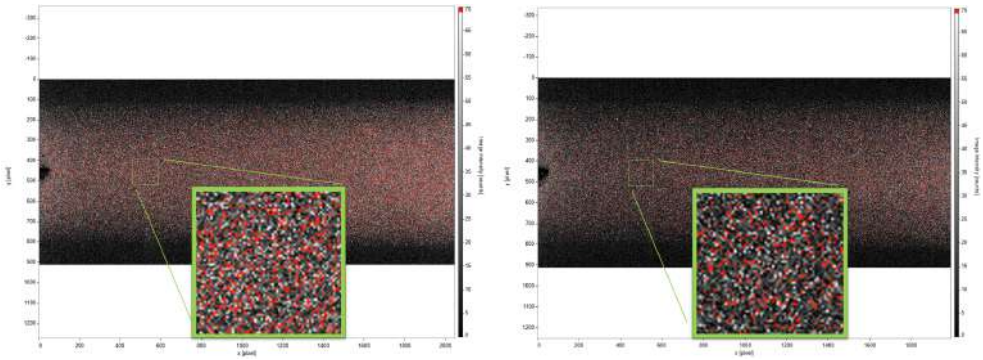

(a) (b)

Figure 3.6: Comparison of particle image density in preprocessed images. (a) Higher-density case with an average particle image density of 0.12 ppp (b) Lower-density case with an average particle image density of 0.08 ppp.

In the next step, the commercial STB implementation is run for all three Reynolds numbers and for two distinct particle image densities in order to assess its robustness with respect to both the flow regime and the seeding level. Figure 3.6 shows two preprocessed images. The panel on the left corresponds to a dataset with a higher particle image density of approximately 0.12 ppp, where the individual particle images start to overlap more frequently. The panel on the right corresponds to a lower particle image density of about 0.08 ppp, for which most particles remain well separated.

Chapter 4

Results

Chapter 4 presents the experimental results obtained from both a selected snow-tracking campaign and the laboratory water-jet measurements and connects them to the theoretical framework introduced in Chapter 1. The first part begins with an assessment of the calibration accuracy of the drone-based imaging system for the snowfall measurements and then focuses on the reconstruction of snowflake Lagrangian trajectories using the IFD algorithm and the corresponding Eulerian fields extracted from this data. The second part presents an analogous analysis for the water-jet tracer particles.

4.1 Drone Calibration Results

In this section, the quality of the snow-tracking calibration is quantified using two complementary error metrics based on the drone-borne light-marker system. An object skew error is first used to evaluate large-scale geometric distortions of the reconstructed marker configuration, followed by an analysis of the image reprojection error that characterises the local consistency between measured image coordinates and the projections of the optimised three-dimensional marker positions.

4.1.1 Object Skew Error

This first error metric quantifies the object skew by comparing the reconstructed three-dimensional marker coordinates to a rigid reference configuration of the drone-mounted calibration object. The latter is obtained by projecting the markers (see Figure 4.1a) with the previously computed projection matrices and the calibration model of Equation 3.1, and then aligning the two configurations via a rigid-body transformation.

Figure 4.1b shows the resulting histogram with the distribution of the skew errors over all calibration frames, in millimetres. The error exhibits a mean value of 1.97 mm with a standard deviation of 1.24 mm, and only a small fraction of samples exceeds 5 mm.

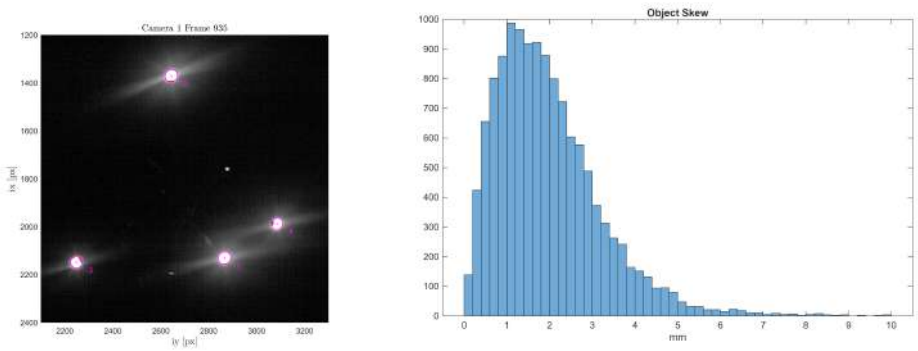

(a) (b)

Figure 4.1: (a) Image of the DJI Matrice 300 light-marker tetrahedron recorded by one camera during the calibration flight. (b) Histogram of the object skew error after optimisation of the drone-based calibration.

4.1.2 Reprojection Error

In this subsection, the calibration accuracy is further evaluated in terms of image reprojection error. For each camera, the reconstructed three-dimensional marker coordinates are projected back onto the image plane using the calibration model of Equation 3.1, and the resulting image coordinates are compared to the measured marker positions. The Euclidean distance between these two sets of points defines the reprojection error in pixels, which is then analysed in terms of its distribution over all cameras and calibration frames.

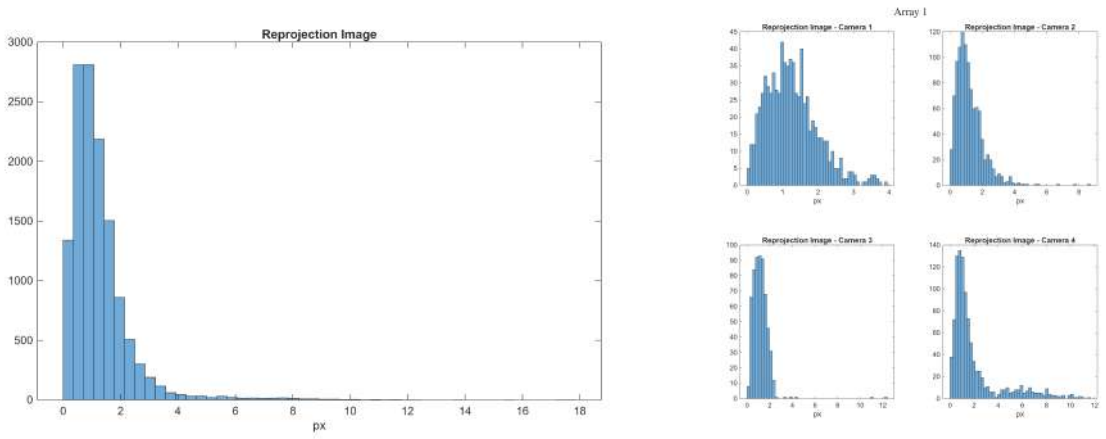

(a) (b)

Figure 4.2: (a) Histogram of the image reprojection error over all cameras and calibration frames, expressed in pixels. (b) Example distributions of the reprojection error for individual cameras within one array.

Figure 4.2 summarises the distribution of the image reprojection error obtained from the drone-based calibration. Panel (a) shows the histogram over all cameras and calibration frames, while panel (b) illustrates the corresponding distributions for individual cameras within one array. The global error exhibits a mean value of 1.38 px with a standard

deviation of 2.77 px, and the vast majority of samples remain below 4 px.

4.1.3 Soloff's Reprojection Error

This part of the analysis assesses the quality of the polynomial Soloff mapping introduced in 3.1.1 by evaluating the associated reprojection errors. For each camera, a regression algorithm is used to identify the polynomial mapping that best fits the calibration data. This mapping is then applied to the three-dimensional marker coordinates to reproject them onto the image plane and compute the corresponding reprojection errors. This procedure has been repeated for polynomial orders ranging from 1 (linear) to 8, and the resulting mean and standard deviation of the reprojection error for camera 2 are reported in Table 4.1.

Table 4.1: Mean and standard deviation of the Soloff reprojection error for camera 2 as a function of the polynomial order. The statistics are computed from 1044 calibration samples.

Order	Mean error [px]	Standard deviation [px]
1	320.12	272.95
2	82.85	69.73
3	20.46	18.25
4	5.65	5.28
5	1.70	1.44
6	0.87	0.58
7	0.96	0.66
8	3.13	2.35

Table 4.1 shows a rapid decrease in both the mean error and the standard deviation as the polynomial order increases until the error reaches a minimum and then starts to rise again. This minimum is attained for the sixth-order polynomial, with a mean reprojection error of 0.87 pixels and a standard deviation of 0.58 pixels, while the corresponding error distribution is shown in Figure 4.3a, approaching error bounds such as those required for STB.

A polynomial function has been fitted for every camera, and the corresponding reprojection errors have been visualised using histograms; an example of these arrays is shown in Figure 4.3b. The associated mean and standard deviation of the reprojection error, as well as the number of calibration samples used for each fit, are summarised in Table 4.2, providing a quantitative measure of the calibration quality for all cameras.

4.2 Snow Tracking Results

In this section, the results obtained from the volumetric snow-tracking experiments are presented. We first assess the calibration quality directly on the reconstructed snowflakes

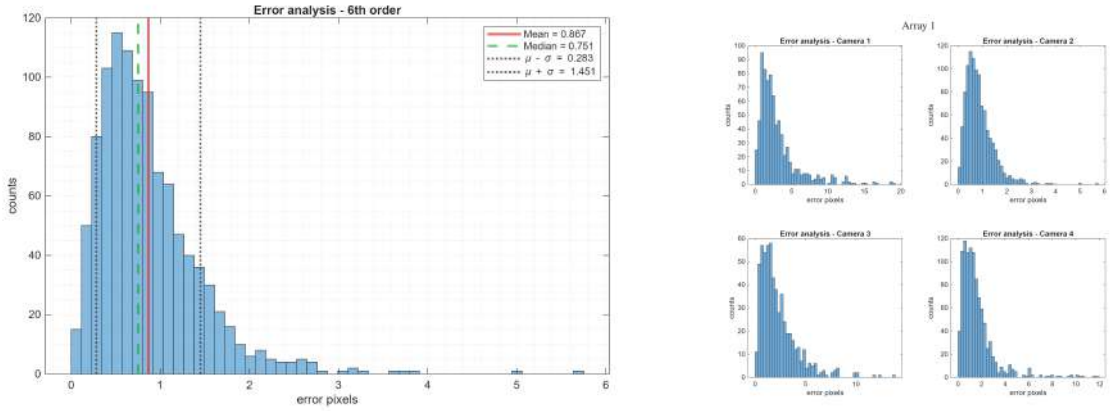

(a) (b)

Figure 4.3: (a) Histogram of the sixth-order Soloff's reprojection error for camera 2. (b) Corresponding sixth-order reprojection error histograms for all cameras in array 1.

Table 4.2: Mean and standard deviation of the Soloff reprojection error for each camera, together with the number of calibration samples used for each fit and the mean base reprojection error. Values highlighted in red indicate poor performance of the Soloff mapping for some cameras.

Camera	Number of samples	Polynomial order	Base mean error [px]	Soloff's mean error [px]	Soloff's standard deviation [px]
1	796	4	1.27	2.30	1.83
2	1044	6	1.22	0.87	0.58
3	600	5	1.19	1.55	1.51
4	1016	7	2.01	1.06	0.94
5	652	5	1.04	0.81	0.50
6	1092	7	1.16	0.95	0.93
7	772	4	1.36	5.69	4.10
8	1024	6	1.28	0.76	0.50
9	628	4	0.97	1.51	1.46
10	664	5	1.19	1.89	1.62
11	820	5	1.06	2.67	2.74
12	812	7	1.33	0.70	0.48
13	604	5	1.24	2.09	1.75
14	876	7	1.18	3.43	2.79
15	656	6	1.08	0.50	0.36
16	1032	6	1.16	0.88	0.81

by analysing their image reprojection error across all cameras. We then examine the raw Lagrangian trajectories provided by the IFD tracking algorithm, before reconstructing the corresponding Eulerian flow fields to extract mean profiles and fluctuation statistics.

4.2.1 Snowflakes Reprojection

As a first step, we evaluate the calibration quality directly on the reconstructed snowflakes by analysing their image reprojection error in each camera. For every particle in a selected frame, the three-dimensional position provided by the IFD tracking algorithm is projected back onto the image plane and compared to the corresponding measured centroid, using both the calibration Equation 3.1 and the Soloff mapping introduced in Section 3.1.1.

Projection-matrix-based reprojection For this reprojection, the camera projection and distortion matrices introduced in Chapter 3 are used to map the three-dimensional particle positions back onto the raw image plane. Figure 4.4 illustrates the pixel-based reprojection on Camera 8 for frame 4. Panel (a) displays the measured and reprojected snowflake centroids on the image plane, showing that the projected positions densely cover the active sensor area and closely follow the observed centroids. Panel (b) shows the corresponding histogram of reprojection error in pixels.

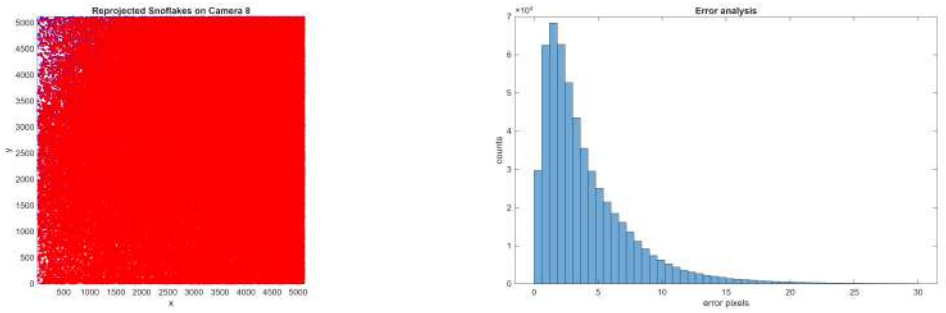

(a) (b)

Figure 4.4: (a) Measured (blue) and camera-matrix-based reprojected (red) snowflake centroids on Camera 8 for frame 4. (b) Distribution of pixel-based reprojection error for Camera 8 in the same frame.

This example is representative of the behaviour observed on the other cameras and provides a visual reference for the summary statistics (median and standard deviation) reported in Table 4.3 for all 16 cameras.

The last column of Table 4.3 reports the average number of occlusions per camera, defined as the mean number of object tracks associated with a single particle detection on the image plane. For four cameras, the corresponding values are not reported because the corresponding arrays were removed from the dataset due to an initialisation issue; these entries are therefore marked as not applicable (-).

Soloff's Reprojection In a second step, we repeat the reprojection analysis using the Soloff polynomial mapping, following the same procedure as described in Section 4.1.3. For each camera, the polynomial order (either fifth or sixth) is selected based on the lowest

Table 4.3: Number of snowflake samples, pixel-based reprojection statistics (mean and standard deviation) and average number of occlusions for each camera in frame 4.

Camera	Number of samples	Mean error [px]	Standard deviation [px]	Average occlusions
1	0	-	-	-
2	0	-	-	-
3	0	-	-	-
4	0	-	-	-
5	302564	5.09	5.52	2.60
6	321371	3.96	3.59	3.03
7	462100	4.09	3.67	3.86
8	549606	4.12	3.78	4.48
9	353267	4.23	4.28	3.03
10	409797	4.89	4.95	3.49
11	384435	3.91	3.81	3.69
12	457801	4.88	4.45	3.75
13	370504	3.82	3.35	4.43
14	499023	3.78	3.64	4.42
15	313772	4.41	4.37	2.68
16	440620	5.29	5.22	3.54

reprojection error and then used to project the reconstructed snowflakes back onto the image plane. Figure 4.5 illustrates the resulting reprojection on Camera 8 for frame 4: panel (a) shows the measured and Soloff-based reprojected centroids on the image plane, while panel (b) presents the corresponding histogram of pixel-based reprojection error.

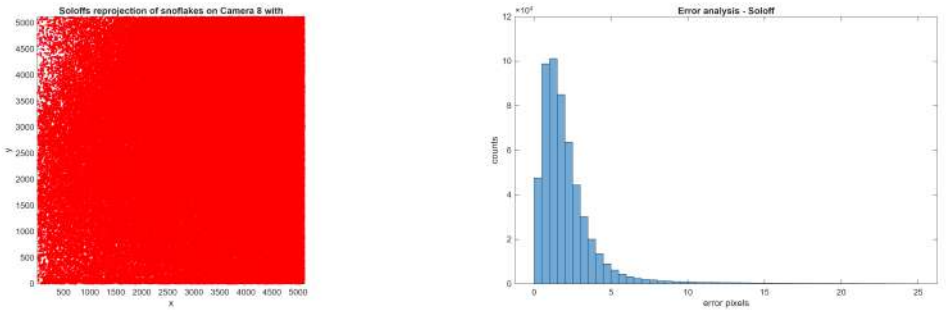

(a) (b)

Figure 4.5: (a) Measured (blue) and Soloff-based reprojected (red) snowflake centroids on Camera 8 for frame 4. (b) Distribution of pixel-based Soloff's reprojection error for Camera 8 in the same frame.

The behaviour observed for Camera 8 is consistent with the results obtained for the other views and serves as a qualitative reference for the statistics reported in Table 4.4. In the following, we therefore summarise the pixel-based Soloff's reprojection error across all 16 cameras in terms of their median and standard deviation.

The first four rows of Table 4.4 contain no values, as Cameras 1–4 do not record any snowflakes for the same reason discussed in the previous paragraph.

Table 4.4: Polynomial order, number of snowflake samples, and pixel-based reprojection statistics (mean and standard deviation) obtained with Soloff's mapping for each camera in the selected frame.

Camera	Polynomial order	Number of samples	Mean error [px]	Standard deviation [px]
1	-	0	-	-
2	-	0	-	-
3	-	0	-	-
4	-	0	-	-
5	5	302564	2.65	2.24
6	6	321371	2.27	2.41
7	5	462100	2.45	2.34
8	6	549606	2.26	2.61
9	5	353267	2.44	1.84
10	6	409797	2.36	2.30
11	5	384435	3.03	2.58
12	5	457801	2.74	2.98
13	6	370504	2.34	1.78
14	6	499023	2.27	2.17
15	5	313772	2.80	2.27
16	6	440620	2.80	3.02

4.2.2 Snow Tracking Data

In the following, we analyse the three-dimensional snow-tracking data obtained using the IFD reconstruction framework, focusing on the spatial distribution of particles and the main properties of the reconstructed trajectories.

Figure 4.6a provides an overview of the reconstructed 10 x 10 x 10 m^3 measurement volume together with the four camera arrays. In total, 294966 tracks of length four frames are shown and coloured by the magnitude of their velocity. Since each track contributes one point per frame, this corresponds to 1179864 particle positions plotted within the volume. Figure 4.6b shows an XZ-slab of 1 m thickness taken through the centre of the measurement volume, which allows for a clearer visualisation of the tracks inside the cube. It is evident that the particle track density within the volume is not uniform: when moving away from the cameras (towards decreasing X and increasing Z), the gaps between neighbouring particles become larger. This qualitative observation is examined more quantitatively in Table 4.5, which reports the number of particles contained in 1 m thick slabs along all three coordinate directions.

The slab statistics in Table 4.5 confirm the qualitative impression. Along the X-direction, slabs closer to the cameras (larger X) contain more particles, although the absolute maximum is not located in the slab at the very edge of the volume. The trend along Z is even clearer: decreasing Z (i.e. moving downward towards the cameras and the light panels) leads to a systematic increase in the number of detected snowflakes. In contrast, the distribution along Y is more symmetric, with the highest particle counts

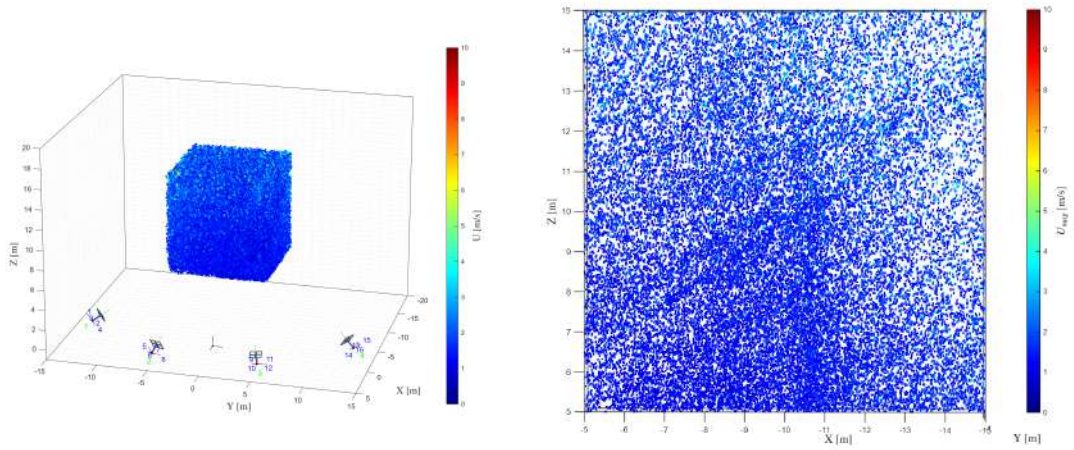

(a) (b)

Figure 4.6: Reconstructed snowflakes coloured by velocity magnitude in the full measurement volume (a) and in a 1 m thick central XZ-slab (b), illustrating the spatial distribution of particles inside the cube.

Table 4.5: Number of reconstructed snowflakes per 1 m thick slab along the three coordinate directions.

X		Y		Z	
Slab range [m]	Snowflakes	Slab range [m]	Snowflakes	Slab range [m]	Snowflakes
(-15; -14)	78401	(-6; -5)	64142	(5; 6)	176164
(-14; -13)	88537	(-5; -4)	76504	(6; 7)	152825
(-13; -12)	99980	(-4; -3)	93673	(7; 8)	138575
(-12; -11)	109436	(-3; -2)	117906	(8; 9)	128175
(-11; -10)	121771	(-2; -1)	141585	(9; 10)	118156
(-10; -9)	131937	(-1; 0)	145149	(10; 11)	110190
(-9; -8)	139420	(0; 1)	145797	(11; 12)	103492
(-8; -7)	142381	(1; 2)	140804	(12; 13)	94214
(-7; -6)	138214	(2; 3)	132764	(13; 14)	82211
(-6; -5)	129775	(3; 4)	121527	(14; 15)	72732

found near the centre of the volume, while one of the two lateral sides (the one closest to the camera array hosting cameras 1–4) exhibits substantially fewer particles than the opposite side.

Another interesting perspective is provided by horizontal slabs extracted at the bottom and at the top of the volume, as shown in Figure 4.7. Panel (a) illustrates the slab at the bottom of the measurement volume, while panel (b) shows the same slab with particles coloured by the Delaunay-based nearest-neighbour distance D, which quantifies the local spacing between snowflakes. The difference in the number of particles and in the typical inter-particle distance within these regions is very pronounced.

A striking feature in Figure 4.7 is the presence of three narrow, streak-like regions where the local particle density is noticeably higher than in the surrounding subvolume. These three lines coincide with the viewing directions of the camera arrays used for the

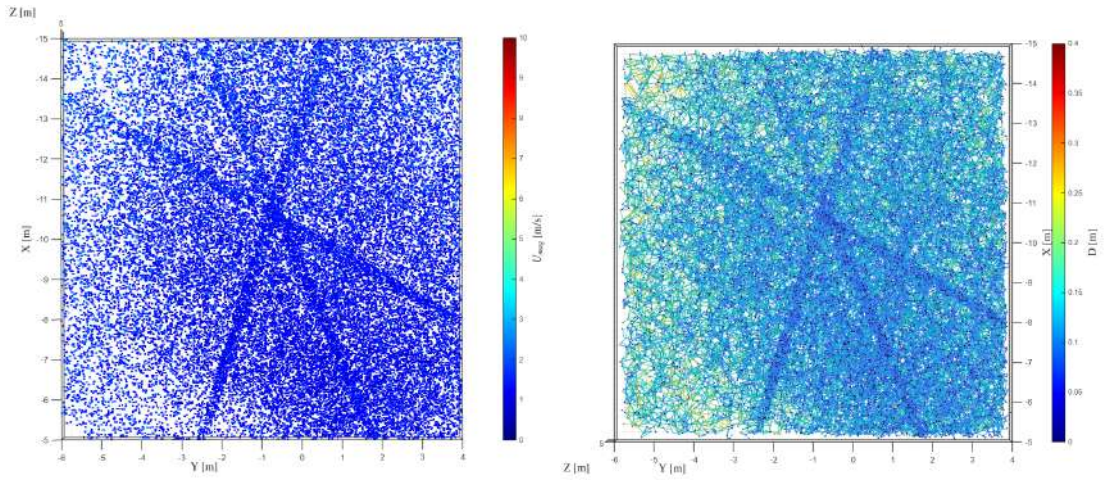

(a) (b)

Figure 4.7: (a) Horizontal slab at the bottom of the measurement volume showing the reconstructed snowflake positions. (b) Same slab, where the particles are coloured by the Delaunay-based nearest-neighbour distance D.

reconstruction of the snowflakes.

Besides the clear difference in particle number, the tracks in the upper slab also appear visually brighter than in the lower slab, suggesting that the corresponding snowflakes move, on average, at higher speeds. To verify this observation, Figure 4.8 shows the mean particle speed in each 1 m thick slab, plotted as a function of its position along the corresponding axis for all three coordinate directions.

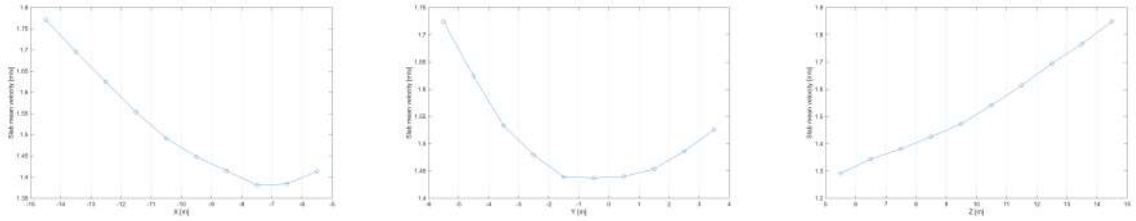

(a) (b) (c)

Figure 4.8: Mean particle speed in 1 m thick slabs as a function of slab position along (a) X, (b) Y, and (c) Z for the snow tracking dataset.

Finally, an overall characterisation of the flow can be obtained by examining the distributions of the three velocity components U, V, and W over the entire measurement volume. Figure 4.9 reports the Probability Density Function (PDF) of U, V, and W computed from all reconstructed tracks, together with a Gaussian fit based on the eleven central bins of each distribution. From these statistics, the mean velocity components within the entire measurement volume are found to be $U_{\text{avg}} = 0.14\ ms^{-1}$, $V_{\text{avg}} = 0.68\ ms^{-1}$ and $W_{\text{avg}} = -1.21\ ms^{-1}$.

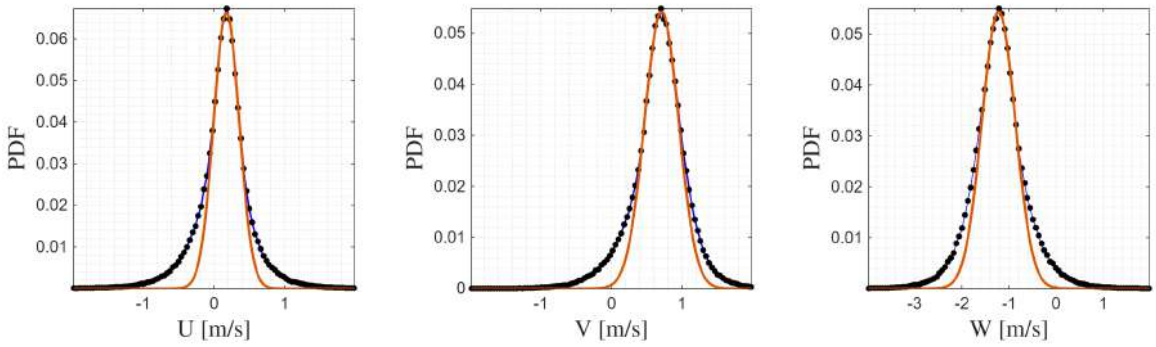

(a) (b) (c)

Figure 4.9: Probability density functions (PDFs) of the three velocity components (a) U, (b) V and (c) W for the snow tracking dataset. The solid lines indicate Gaussian fits obtained from the eleven central bins of each distribution.

4.2.3 Extracting Field Descriptions

In this section, the Lagrangian snowflake trajectories are used to construct Eulerian velocity fields and to quantify the corresponding flow fluctuations within the measurement volume. The Eulerian velocity fields are reconstructed by locally fitting a smooth spatio-temporal model to the Lagrangian trajectories on a regular grid and evaluating the resulting displacement field and its derivatives at each grid point.

Figure 4.10a provides an overview of the entire measurement volume, displaying velocity vectors at each grid point that have been time-averaged over the four frames and coloured by their magnitude to visualise the three-dimensional structure of the mean flow.

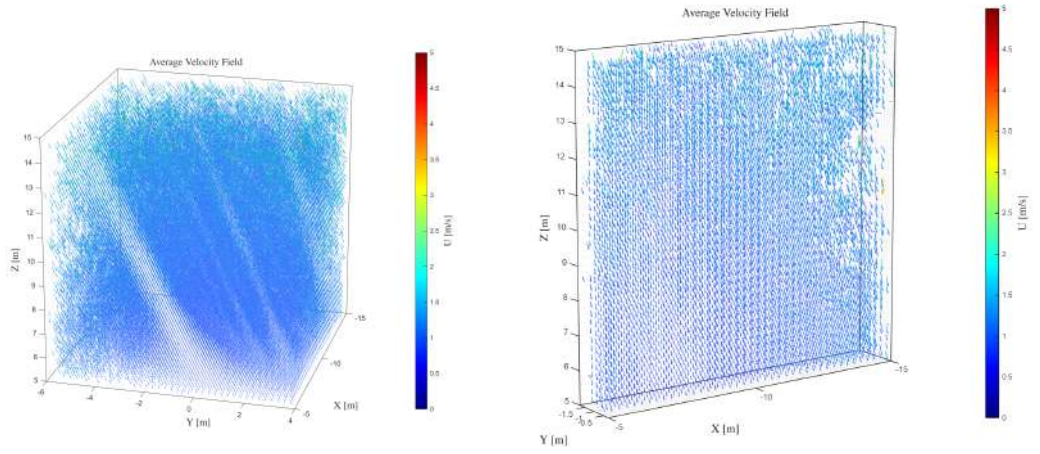

(a) (b)

Figure 4.10: Time-averaged Eulerian velocity field reconstructed from the snowflake trajectories. (a) Overview of the entire measurement volume, showing velocity vectors at each grid point coloured by their magnitude. (b) One-meter-thick slab extracted from the centre of the volume along the Y-direction.

.

Figure 4.10b shows a 1-meter-thick slab extracted from the centre of the measurement volume along the Y-direction. In this view, the velocity vectors are more clearly visible: the region of the slab closer to the cameras exhibits an almost uniform mean flow, whereas towards the opposite side, both the magnitude and direction of the velocity show noticeable variations.

In addition to the mean velocity field, the turbulent velocity fluctuations offer a measure of the local mixing intensity and quantify how strongly the instantaneous snowflake motion departs from the average flow. For each grid point, the instantaneous velocity components are first decomposed into a mean and a fluctuating part:

$$U = \langle U \rangle_{xyz} + u', \qquad V = \langle V \rangle_{xyz} + v', \qquad W = \langle W \rangle_{xyz} + w' \tag{4.1}$$

The fluctuating components u', v' and w' were obtained by subtracting the spatially averaged velocity from the instantaneous velocity at each grid point. Because the particle trajectories span only four frames, a statistically meaningful time average cannot be computed. Therefore, the mean velocity was estimated as the spatial average over the entire measurement volume, under the assumption that the mean flow is approximately uniform within the observed domain. The resulting fluctuation field can be visualised directly in physical space. Figure 4.11a presents the three-dimensional distribution of the fluctuation vectors for frame 4, coloured according to their magnitude to highlight regions where the instantaneous velocity departs most strongly from the spatially averaged flow.

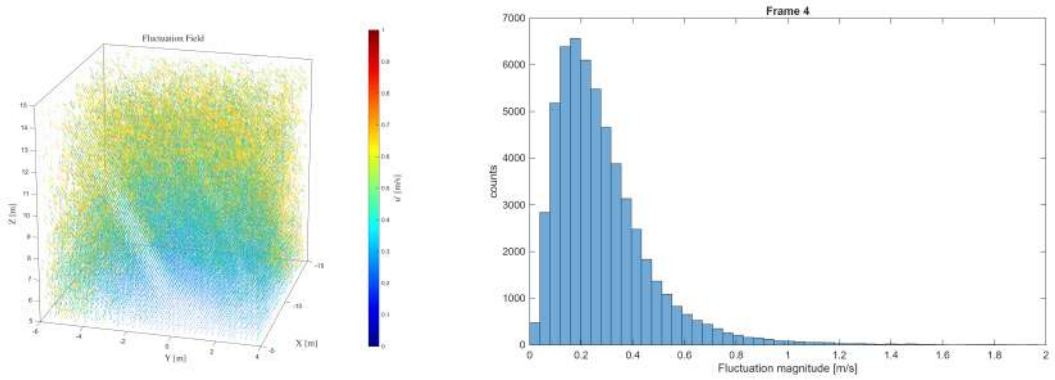

(a) (b)

Figure 4.11: (a) Three-dimensional fluctuation field within the measurement volume, with vectors coloured by their instantaneous magnitude. Note that the colour scale differs from the one used for the mean velocity fields in Figure 4.10. (b) Histogram of the fluctuation magnitude for frame 4, illustrating the distribution of turbulent velocities sampled by the snowflakes.

Figure 4.11b reports the histogram of the fluctuation magnitude for frame 4, providing a compact statistical view of the range of turbulent velocities sampled by the snowflakes.

4.3 Water Jet Tracking Results

This section presents the results of the laboratory water-jet experiment. First, the Lagrangian particle tracks obtained with the commercial and in-house STB algorithms are analysed, followed by the corresponding Eulerian fields reconstructed with FF3 and the resulting mean velocity and turbulence statistics.

4.3.1 STB Data from DaVis

In this section, the Lagrangian particle tracks obtained with the commercial STB algorithm for the turbulent water-jet experiment are presented for the three investigated Reynolds numbers.

Figure 4.12 illustrates representative visualisations of the reconstructed trajectories. Panels (a), (b), and (c) show the full measurement volume for $Re = 2000$, 4000, and 11000, respectively, whereas panels (d), (e), and (f) display the corresponding 10 mm-deep central plane sections used to highlight the internal jet structure; in all panels, the jet nozzle is located on the right-hand side of the measurement volume, and the particles therefore move from right to left. It should be noted that the colour scale is based on the absolute velocity of each particle, but the colour range is not the same for all Reynolds numbers, as the jet-exit velocity varies from 0.2 ms^{-1} at $Re = 2000$ to 1.2 ms^{-1} at $Re = 11000$.

The particle image density is not constant across the different cases: panels (a) and (d) correspond to a density of approximately 0.08 ppp, panels (b) and (e) to about 0.12 ppp, and panels (c) and (f) to about 0.10 ppp. This effect is clearly visible in Figure 4.12, where the images with higher particle image density appear noticeably darker due to the larger number of particles compared to the lower-density cases.

Finally, some high-velocity tracks can be observed on the left-hand side of the measurement volume, particularly in the high-density cases. These non-physical trajectories are so-called ghost tracks, and their origin and impact will be further analysed in Chapter 5.

The DaVis software provides access to the residual images generated by the STB algorithm, which can be used to assess its performance. These images (see Figure 4.13b) are obtained by reprojecting the reconstructed three-dimensional particle positions onto the image plane and subtracting the corresponding synthetic particle images from the original recording (Figure 4.13a). In an ideal reconstruction, the residual intensities would be close to zero everywhere; any bright structures in the residual images, therefore, correspond to particles that are either incorrectly tracked or not properly removed by the algorithm.

From Figure 4.13b, corresponding to the $Re = 11000$ case, it can be seen that the residual amplitudes remain low over most of the FOV. Only the regions near the jet nozzle

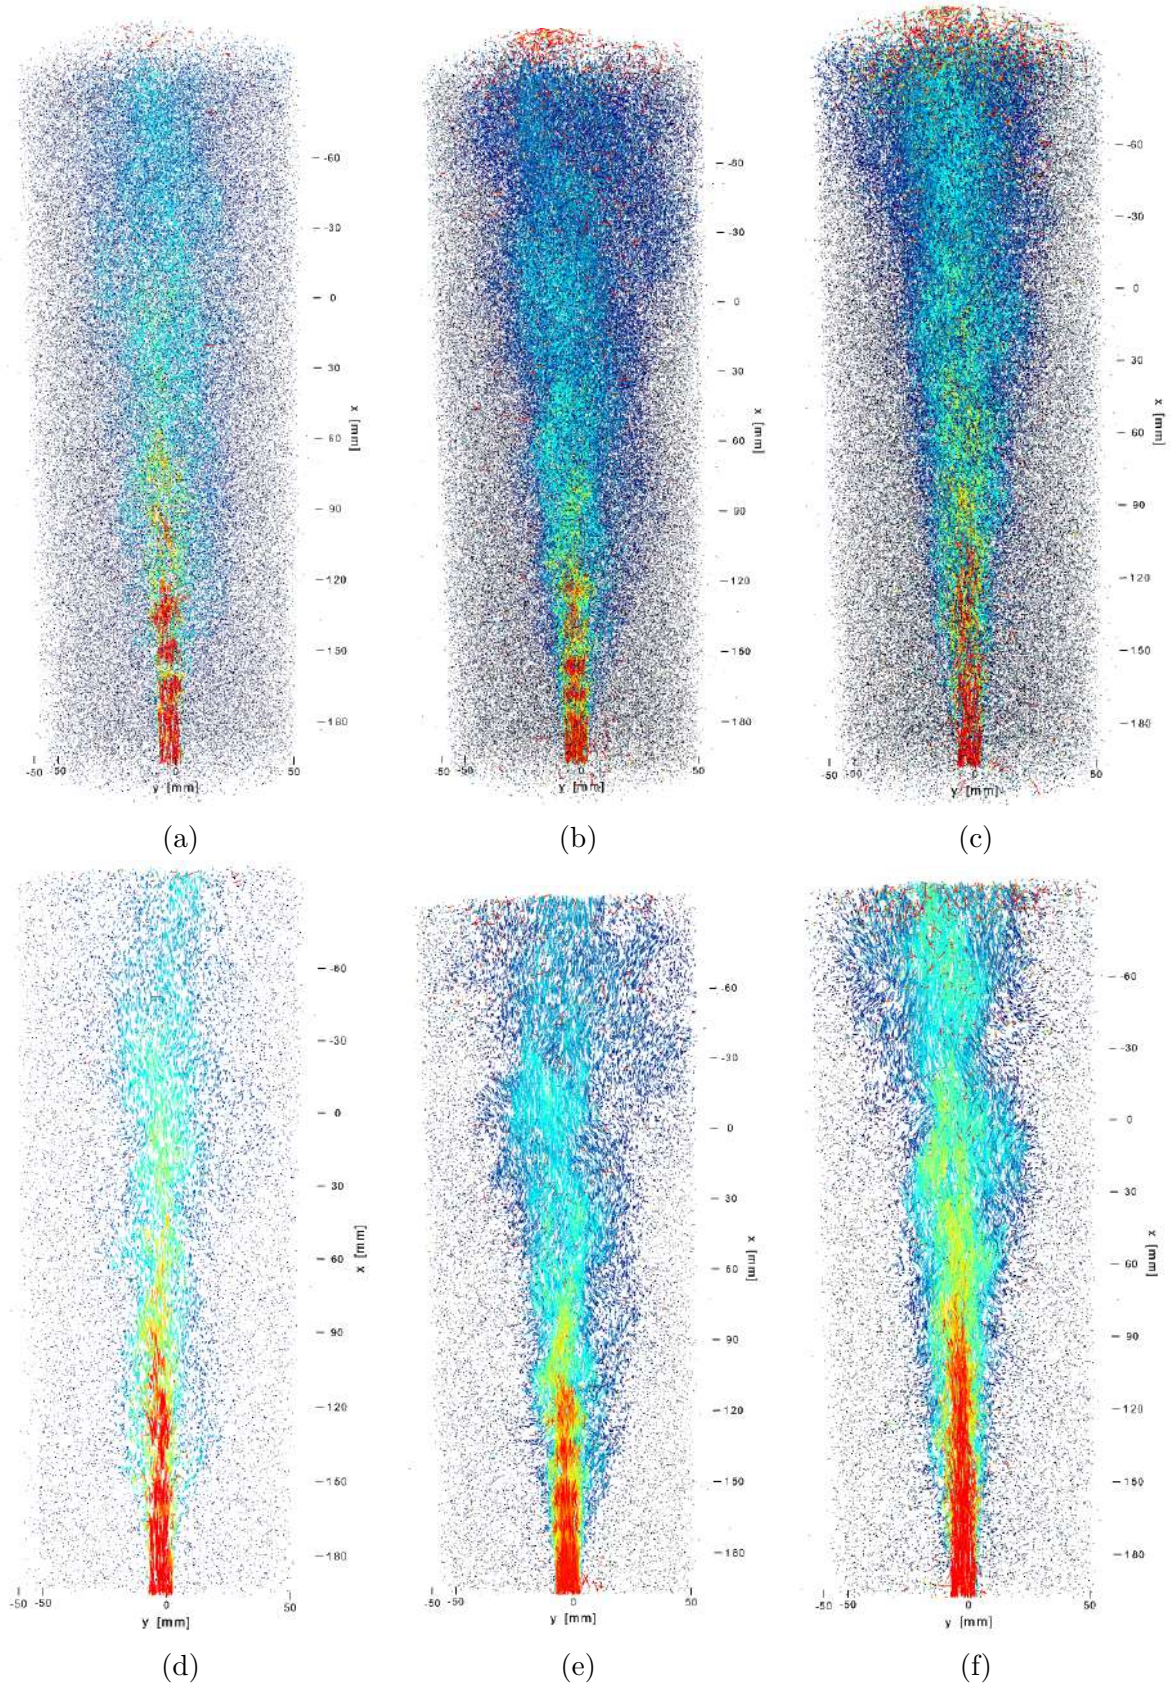

Figure 4.12: Lagrangian particle tracks obtained with STB. Panels (a), (b), and (c) show the full measurement volume for Reynolds numbers $Re = 2000$, 4000, and 11000, respectively, using a track length of 15 frames. Panels (d), (e), and (f) show the corresponding central plane sections, defined by $-5 \leq z \leq 5$ mm, using a track length of 29 frames.

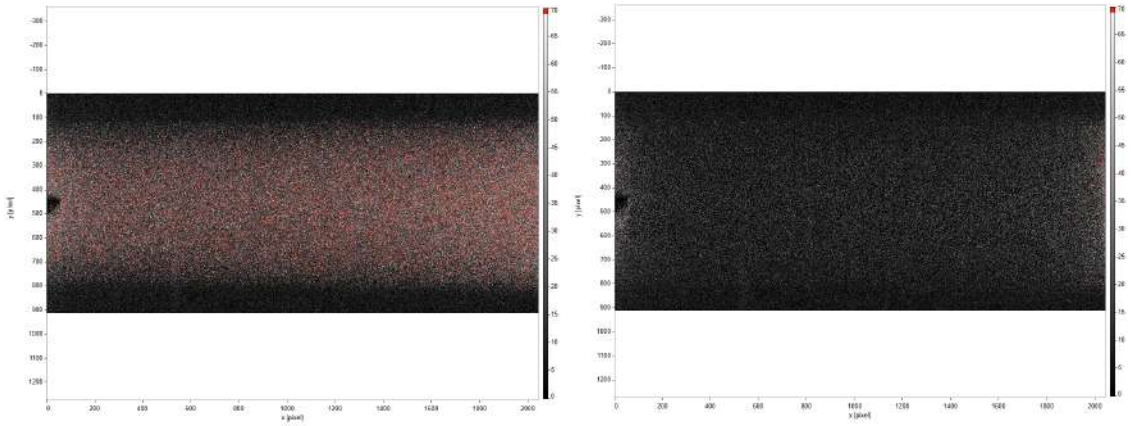

(a) (b)

Figure 4.13: (a) Preprocessed particle image for $Re = 11000$. (b) Residual image of the same frame after STB processing.

on the left-hand side of the panel and towards the end of the measurement volume on the right-hand side exhibit relatively high-intensity peaks.

Beyond the qualitative appearance of the trajectories, it is useful to characterise the STB performance in terms of particle and track statistics, namely the number of detected particles, total tracks, and newly created tracks as a function of time. The DaVis software also provides access to basic statistics for STB processing, such as the number of detected particles, the total number of tracks, and the number of newly created tracks at each time step. For the intermediate case $Re = 4000$, these statistics are shown in Figure 4.14

The phases of the STB algorithm described in Section 1.4 can be clearly identified in these graphs. The convergence phase is characterised by a rapid increase in the total number of tracks and a high rate of newly created tracks, followed by a quasi-steady regime in which both quantities fluctuate around approximately constant values.

Panel (a) shows the statistics obtained from a single-pass STB reconstruction. In this case, the converged state is reached after slightly fewer than 200 time steps, when the total number of tracks saturates, and the number of newly created tracks drops to a low, approximately constant level. Panel (b) instead reports the results of a multi-pass reconstruction in which the STB algorithm is applied several times to the same dataset using different time increments.

The main idea is to enable the tracking of slow-moving particles, typically those entrained towards the jet, whose frame-to-frame displacement is too small to be reliably linked in a standard pass, especially in the presence of overlapping particle images that can remain superposed for many frames. These additional passes can also be performed backwards in time, so that the convergence phase occurs near the end of the sequence and the initial frames benefit from an already well-populated set of tracks [24]. This effect is clearly visible in panel (b), where the converged state is reached in fewer than 100 time

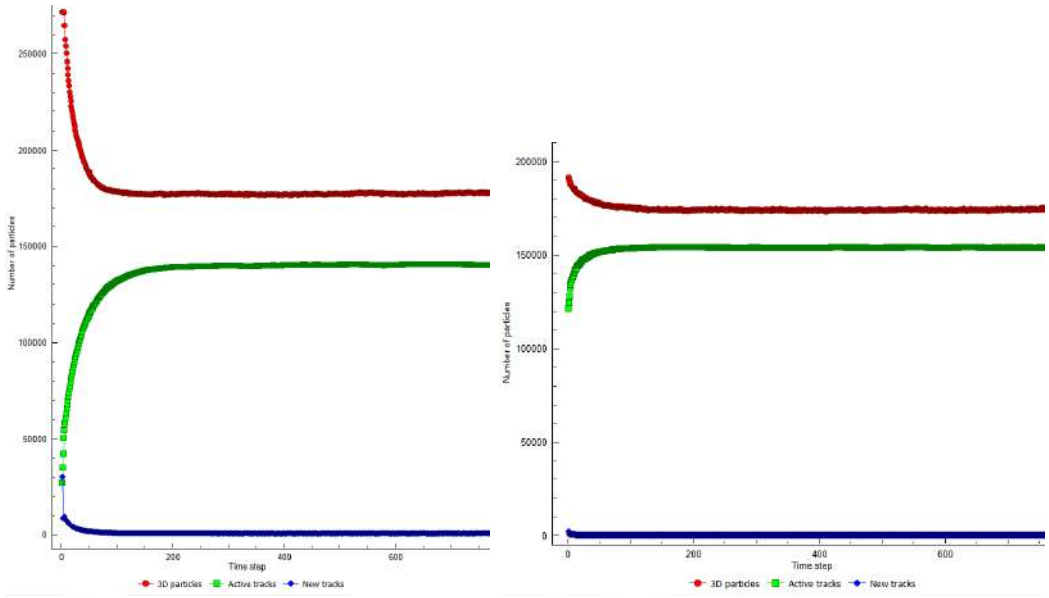

(a) (b)

Figure 4.14: STB statistics for $Re = 4000$. The number of detected particles (red), total tracks (green), and new tracks per time step (blue) are shown as a function of time step index. Panel (a) corresponds to a single-pass reconstruction, panel (b) to the multi-pass reconstruction.

steps, significantly earlier than in the single-pass case.

4.3.2 Fine Scale Reconstruction Data

In this section, the Lagrangian particle tracks are further processed using the Fine Scale Reconstruction (FSR) method to obtain three-dimensional Eulerian fields for the turbulent round jet. Analogously to the FF3 data assimilation framework, FSR projects the scattered Lagrangian velocities onto a smooth, nearly divergence-free Eulerian field on a regular grid [35]; however, it is implemented as an integrated post-processing module within DaVis and is generally less flexible and slightly less performant than FlowFit in terms of reconstruction quality and computational efficiency for large datasets.

Figure 4.15 shows instantaneous velocity fields reconstructed with FSR for the three investigated Reynolds numbers. In all cases, the jet issues from the nozzle (right of the panels) and develops into the surrounding quiescent fluid with a pronounced high-velocity core along the centerline and decreasing velocities towards the jet edges. As the Reynolds number increases, the reconstructed fields exhibit higher peak velocities and increasingly fine-scale turbulent structures; at the same time, the overall jet shape and spreading remain qualitatively similar across all three cases.

The time-averaged field obtained by FSR for the lowest Reynolds number is reported in Figure 4.15d. The mean velocity distribution displays a narrow, high-speed jet close to the nozzle exit, followed by a gradual centerline decay and radial broadening further

downstream. In some regions, small residual fluctuations are still visible where a smoother mean profile would be expected. This indicates that the averaging window and/or the underlying reconstruction parameters are not sufficient to fully suppress instantaneous variability, and the mean field should therefore be interpreted with some caution.

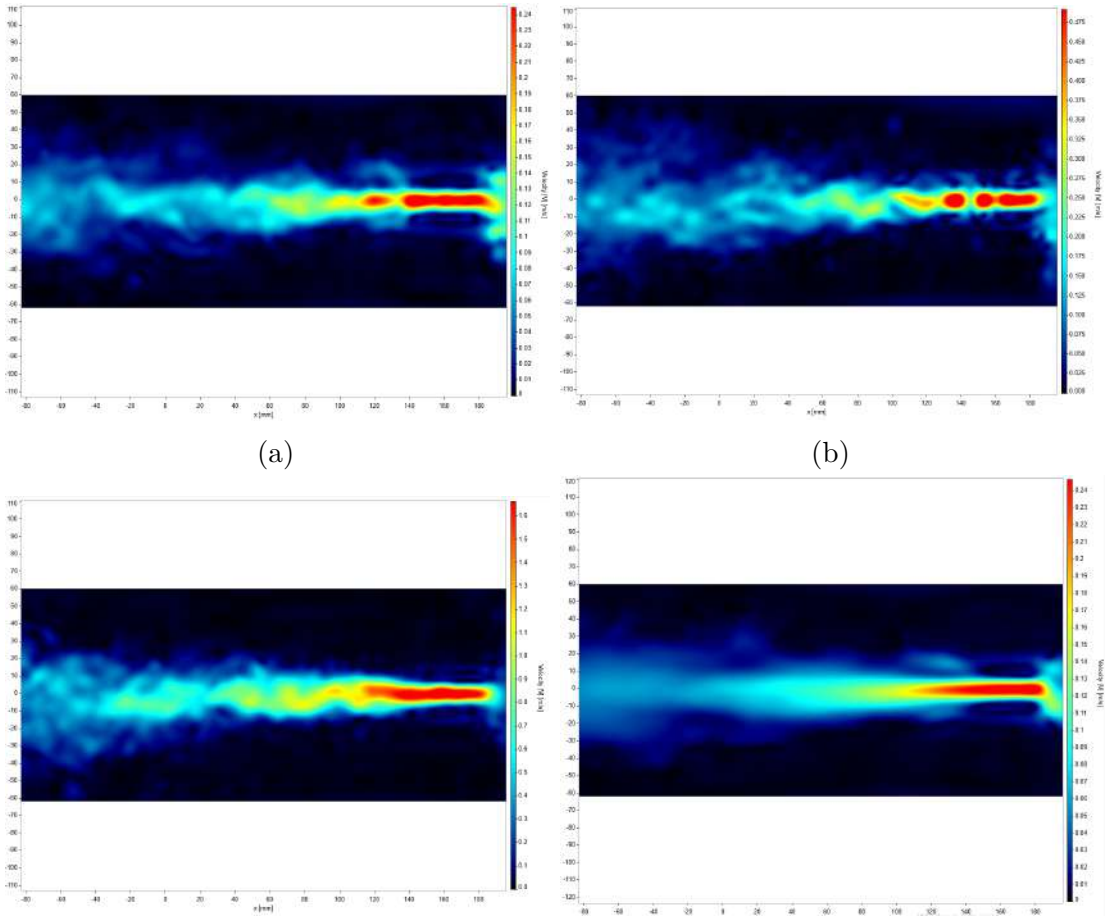

(a) (b)

(c) (d)

Figure 4.15: Eulerian velocity fields reconstructed with the FineScaleReconstruction method. Panels (a), (b) and (c) show instantaneous fields for $Re = 2000$, $Re = 4000$ and $Re = 11000$, respectively, while panel (d) shows the time-averaged field for $Re = 2000$ over 2000 frames.

Chapter 5

Discussion

This chapter discusses the main findings obtained from the field measurements of natural snowfall and the laboratory experiments on the turbulent water jet. Building on the results presented in Chapter 4, the analysis first focuses on a comparative assessment of the two experimental approaches, highlighting their respective strengths and limitations in measurement quality and achievable flow diagnostics. The discussion then turns to the underlying flow physics, examining how the reconstructed boundary-layer snow trajectories and the self-similar jet profiles relate to classical turbulence theory and prior studies.

5.1 Comparison between Field and Lab Experiments

The snow and water-jet experiments represent two extreme points in terms of measurement strategy: a large, complex outdoor setup with limited control over ambient conditions and a high fidelity laboratory configuration where most parameters can be carefully controlled. This section does not re-describe the setups in detail but instead compares how these contrasting conditions impact the underlying measurement quality. In particular, the following subsections address calibration accuracy, spatial particle coverage, and reconstruction artefacts such as ghost particles and occlusions, which are key to interpreting the results presented in Chapter 4.

5.1.1 Calibration Accuracy and Measurement Uncertainty

A key difference between the two measurement configurations lies in the achievable calibration accuracy. Accurate calibration is essential for robust particle tracking because any systematic bias in the camera–volume mapping directly propagates into errors in the reconstructed particle positions and velocities. This is particularly critical for high-density Lagrangian methods such as STB, which rely on sub-pixel accurate reprojection of synthetic particle images onto the camera sensors; if the calibration is not sufficiently precise and spatially homogeneous, the iterative “shaking” and track prediction steps can no longer

reliably distinguish true particles from ghost images.

In the laboratory water jet experiment, this requirement is largely met. The combination of a freely moved calibration target, VSC, and in particular its OTF reconstruction described in Section 3.2 yields a reprojection error on the order of a few hundredths of a pixel across the whole measurement volume. This sub-pixel accuracy is close to what can be considered ideal for high-density STB processing and therefore provides a natural reference level for calibration quality.

Achieving a comparable level of accuracy in large-scale outdoor measurements is considerably more challenging. In the snow-tracking experiment, the cameras span a much larger baseline and observe a volume that is several orders of magnitude larger under varying atmospheric conditions and illumination. Despite these constraints, the calibration results obtained for the snow-tracking campaign are better than might be expected for such a large outdoor setup. The direct pinhole-based calibration described by Equation 3.1 already yields reprojection errors of a few pixels for both the drone markers and the snowflake images. These errors are further reduced when the Soloff-based mapping is applied, reaching sub-pixel accuracy for some cameras in the case of the drone markers and limiting the snowflake reprojection error to only a few pixels across a vast volume.

Overall, the gap in calibration accuracy between the two experiments remains substantial, with the water-jet facility achieving errors that are one to two orders of magnitude smaller in pixel space. However, this difference must be interpreted in the context of the vastly different measurement scales: the snow-tracking system covers a 10 x 10 x 10 m^3 volume, and the typical spacing between neighbouring snowflakes is much larger, around one order of magnitude, than the separation between tracer particles in the laboratory jet.

5.1.2 Spatial Distribution of Particles

Another way to assess the quality of the measurements is to examine the spatial distribution of particles, i.e., the particle number density throughout the measurement volume. In the water-jet experiment, the seeding conditions were well controlled and kept approximately constant within the illuminated region during the acquisition. This is reflected in the reconstructed fields, where the particle density appears uniform across the volume and the total number of tracked particles remains nearly constant during the converged phase of the processing. Together with the underlying conservation principles imposed in the data assimilation, this indicates that the reconstructed tracer distribution is physically consistent and can be considered a faithful representation of the actual flow.

For the snow-tracking experiment, the situation is different. In this case, the inflow and outflow of snowflakes through the measurement cube cannot be controlled, so the number of particles inside the volume naturally fluctuates over time as the atmospheric conditions evolve. Nevertheless, under steady snowfall conditions, it is reasonable to assume that

the mean particle concentration in the air is approximately uniform over the scale of the measurement domain, so that, in the absence of measurement biases, the reconstructed number density should not exhibit strong large-scale gradients.

In practice, the reconstructed particle distribution does not exhibit this approximate uniformity. Instead, the horizontal slabs in Figure 4.7 reveal pronounced bands of increased number density, where snowflakes are more numerous and closer to each other. These bands align closely with the main lines of sight of the camera arrays, indicating that the observed variations in particle density are primarily caused by the imaging geometry and reconstruction biases rather than by genuine inhomogeneities in the snowfall. This reconstruction bias is particularly problematic because it is accompanied by systematic variations in the measured particle velocities. The slab-wise statistics in Figure 5.1 show that regions with higher reconstructed number density also tend to exhibit lower mean normalised speeds, whereas sparser regions appear to host faster snowflakes.

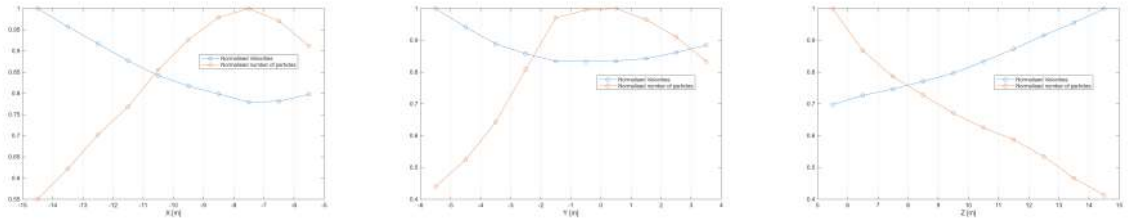

(a) (b) (c)

Figure 5.1: Mean normalised particle speed (blue) and number of snowflakes (orange) in 1 m thick slabs as a function of slab position along (a) X, (b) Y, and (c) Z for the snow tracking dataset.

A plausible explanation for the observed anti-correlation between particle number density and mean normalised speed is that the regions with fewer detected snowflakes coincide with the portions of the volume that are most difficult to image reliably. In particular, these regions tend to lie farther away from the active camera arrays and outside their main lines of sight, which results in reduced effective sensitivity and increased reconstruction bias. The situation is further exacerbated by the fact that one of the four arrays is not used in the present reconstruction, thereby amplifying the anisotropy of the imaging geometry and leaving larger sub-volumes that are only weakly constrained by the available views.

In such poorly constrained regions, the tracking framework is more likely to miss small, dim, or partially occluded flakes, while still detecting larger and brighter particles that typically have higher terminal fall velocities [36]. As a consequence, the local ensemble of reconstructed trajectories becomes biased towards heavier and faster snowflakes, which artificially increases the slab-averaged speed and produces the apparent trend that the flow "seems faster" where fewer particles are detected. This interpretation supports the view that the slab-wise velocity variations in Figure 5.1 primarily reflect a measurement and

reconstruction bias, rather than genuine large-scale inhomogeneities of the atmospheric flow.

5.1.3 Line of Sight Ambiguities and Occlusions

In the snow-tracking configuration, a major source of uncertainty arises from line of sight ambiguities in the multi-camera reconstruction. In this context, a line of sight is the straight ray that connects a camera centre to a point in the physical domain (here, a particle). The intersection between this ray and the image plane defines the pixel location at which that particle appears in the recorded image, as illustrated schematically in Figure 5.2a. A second camera, placed at a different location and orientation, defines its own line of sight to the same particle in three-dimensional space. The particle position is then obtained as the point where the viewing rays from the two cameras intersect (or come closest to intersect), which is the basic principle of triangulation.

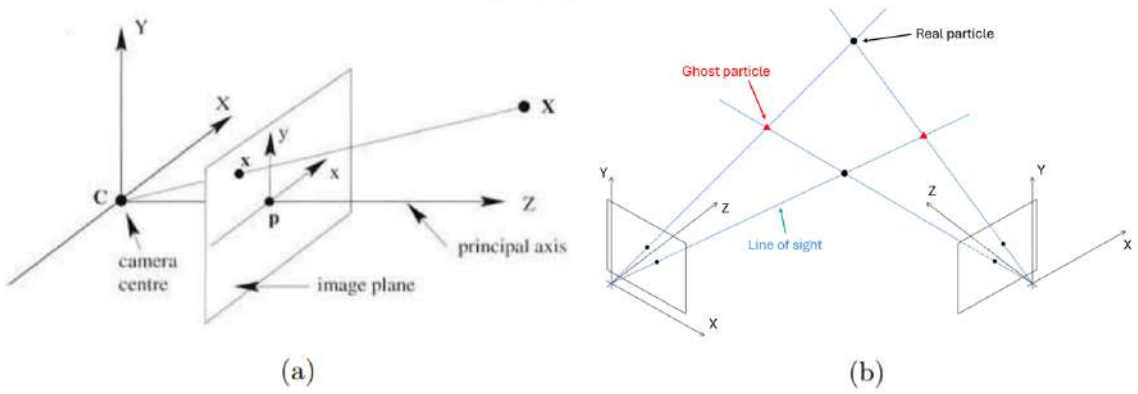

Figure 5.2: (a) Schematic representation of the line of sight geometry, reported from [11]. (b) Example of an incorrectly reconstructed particle due to line of sight ambiguities.

A problem arising from this reconstruction is illustrated in Figure 5.2b, shown here for the case of two cameras. For a certain configuration of a pair of particles in object space, the four lines of sight (two per camera) can intersect at four distinct locations: two intersections coincide with the true particle positions, while the other two occur at positions where no particle is actually present, thereby creating ghost particles. In a volume with a high particle number, such ambiguous configurations are expected to occur frequently. The most straightforward way to reduce this issue is to add more cameras so that a valid reconstruction requires the intersection of more than two lines of sight. This is why four views are commonly employed in volumetric LPT systems, offering enough redundancy.

A distinctive feature of these ghost particles is that they are reconstructed along the line of sight of real particles, so that both the ghost and the physical particle are projected onto (approximately) the same location on the image plane. This effect is referred to as

occlusion. Table 4.3 reports the mean number of occlusions observed in the cameras for a single frame. This metric quantifies, on average, how many reconstructed particles lie along the line of sight of each detected image feature, and it is between three and four for most cameras. While it is plausible that some viewing rays genuinely intersect several snowflakes at different depths, the fact that the mean value is significantly larger than unity strongly suggests that a considerable fraction of these assignments corresponds to ghost particles arising from the line-of-sight ambiguities discussed above. As a result, the high average number of occlusions per detected image particle introduces a systematic bias in the reconstructed snow distribution and associated statistics.

This unusually high average number of occlusions is partly a consequence of the limited redundancy of the present imaging system, where only three camera arrays are used for the volumetric reconstruction. In addition, the particle tracks are evaluated over just four consecutive frames, which strongly restricts the amount of temporal information available to resolve ambiguous matches. Together, the reduced number of views and the very short track length make it much easier for ghost particles generated by line-of-sight ambiguities to survive the reconstruction and appear as additional occlusions in the snow-tracking data. In contrast, the STB processing used for the water-jet data benefits from both a larger number of camera views and much longer particle tracks, which provide enough geometric and temporal redundancy to resolve most line-of-sight ambiguities. This explains why ghost particles in the water-jet reconstructions can be reliably identified and filtered based on their kinematics, whereas in the snow-tracking configuration a comparable level of ambiguity remains in the data.

5.1.4 Ghost Particles in the Water Jet Data

Even the STB reconstruction is not immune to ghost particles. As shown in Figure 4.12, at the highest seeding densities and towards the downstream end of the measurement volume, some reconstructed tracks exhibit unrealistically large velocities that cannot be explained by the jet dynamics. In contrast to the ghosts arising in the snow reconstruction, these spurious particles in the water-jet data are relatively easy to detect. Since their apparent velocities exceed the physically plausible range by a large margin, they can be effectively removed by applying a simple threshold on particle speed, thereby allowing a straightforward statistical quantification of their occurrence.

Figure 5.3 shows the PDF of the particle speed in a single frame, evaluated from particles located in the last 20 mm of the measurement volume in the downstream region. When plotted on a linear scale, the PDF does not exhibit any clear signature of high-speed particles because the fraction of ghost particles remains very small. In contrast, the insets plotted on a semi-logarithmic scale reveal a distinct contribution from these rare, high-speed events. For the $Re = 2000$ case, where the particle image density is relatively

low, the PDF is dominated by a strong peak at small velocities, and the tail at higher speeds is almost negligible. At higher seeding densities (e.g. $Re = 11000$), the initial peak becomes weaker, and the high-speed tail gains relative weight, indicating a more significant contribution of spurious, non-physical velocities associated with ghost particles.

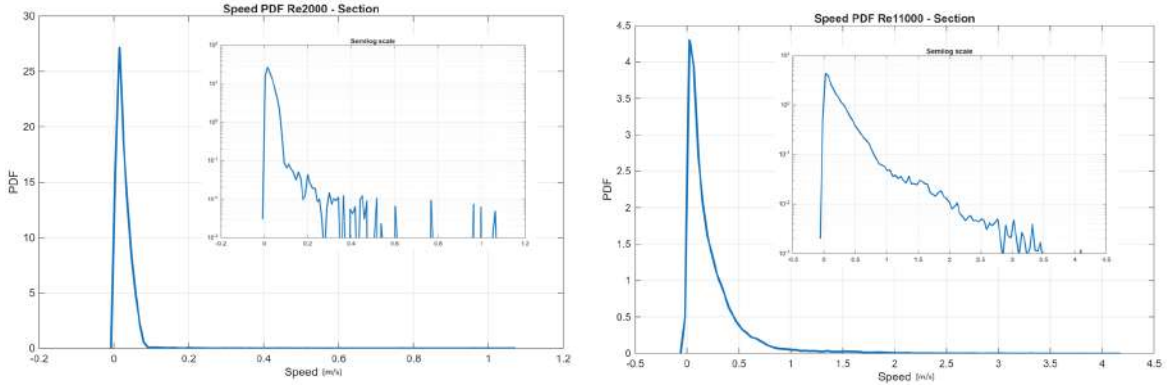

(a) (b)

Figure 5.3: Probability density functions of the particle speed in the water-jet experiment, evaluated from particles located in the last 20 mm of the measurement volume for (a) $Re = 2000$ and (b) $Re = 11000$; the insets show the same distributions on a semi-logarithmic scale.

In the present work, to quantify the occurrence of ghost particles, the threshold is defined based on the jet-exit velocity U_J. Particles with an apparent speed larger than U_J in the last 20 mm of the measurement volume are considered dynamically inconsistent with the flow and are therefore classified as ghosts. For the low particle density case at $Re = 2000$, only 0.37% of the particles exceed the threshold; this fraction increases to 1.35% for $Re = 4000$ and 1.58% for $Re = 11000$, indicating that the occurrence of such high-speed outliers likely depends not only on the particle image density but also on the increasing flow velocity and velocity variation.

5.2 Analysis of the Flows Physics

The previous section focused on the measurement aspects of the two experiments, highlighting differences in calibration accuracy, particle distribution, and reconstruction analysis. In the following, the discussion shifts to the underlying flow physics that can be extracted from the reconstructed data. First, the snow measurements are interpreted in the framework of an atmospheric turbulent boundary layer, using the Lagrangian and Eulerian statistics from Chapter 4 to assess how well the observed flow conforms to canonical boundary layer behaviour. Subsequently, the water-jet experiment is analysed in terms of its mean velocity profile and centerline velocity and compared to the self-similar round-jet scalings introduced in Chapter 1.

5.2.1 Snow Boundary Layer

In this section, we interpret the reconstructed snow particle velocities in the turbulent boundary layer by discussing the organisation of the mean flow, relating it to the theoretical framework introduced in Chapter 1, and comparing the results with previous work [10, 32] to assess their physical consistency. Before discussing the flow behaviour, it is necessary to identify in which region of the ABL our measurements take place by transforming the vertical coordinate Z into the dimensionless wall distance.

$$z^+ = \frac{Z u_\tau}{\nu} \tag{5.1}$$

Where Z is the distance from the ground, u_τ is the friction velocity, and ν is the kinematic viscosity of air, here assumed to be $1.5 \cdot 10^{-5} m^2 s^{-1}$. Due to the very small kinematic viscosity, the wall distance in viscous units increases rapidly with height, so that even moderate geometric distances from the surface correspond to large z^+. For the measurement cube considered here, the lower boundary at $Z \approx 2m$, already lies well above the conventional threshold $z^+ > 50$, indicating that the measurements are located within the logarithmic region of the boundary layer rather than in the viscous sublayer or buffer layer.

Figure 5.4 shows the law of the wall for a turbulent boundary layer in (a) a linear and (b) a semi-logarithmic representation. The blue, green, and black lines correspond to the solutions of the BLE in the viscous sublayer, buffer layer, and logarithmic region, respectively, while the red symbols indicate the present snow–particle measurements expressed in wall units.

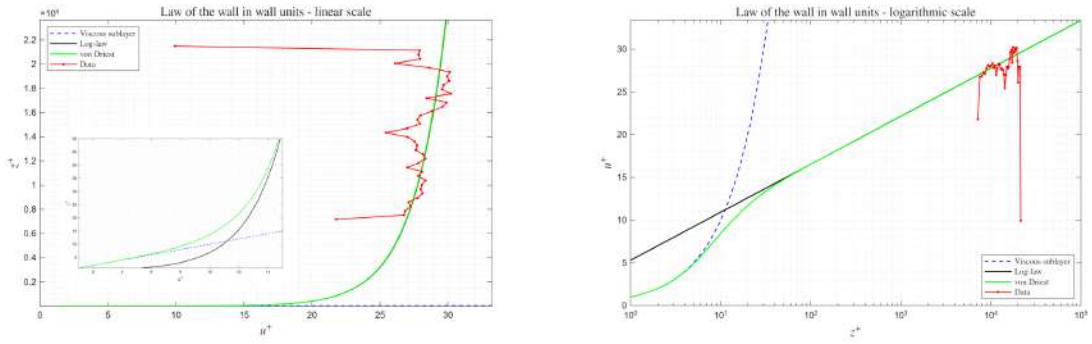

(a) (b)

Figure 5.4: Law-of-the-wall representation of the mean horizontal snow–particle velocity profile in dimensionless viscous units. (a) Linear scale, showing the measured data (red) together with the viscous sublayer, logarithmic law, and von Driest interpolation. (b) Logarithmic representation of the same profile

To obtain the wall-unit representation of the snow-particle velocity and height, the logarithmic law in Equation 1.2 is fitted to the data over the accessible height range in order to estimate an effective friction velocity u_τ, such that the resulting data points

approximately follow the theoretical behaviour. The physical relevance of this procedure, however, depends critically on the resulting value of u_τ, which in our case is $0.0215ms^{-1}$. Looking at previous similar works, [10] recovered a friction velocity of $u_\tau \approx 0.3ms^{-1}$, while [32] reported $u_\tau \approx 0.24ms^{-1}$; the present estimate is therefore almost one order of magnitude smaller than these reference values. This difference can be explained by the much weaker wind speed in the present configuration: the average horizontal velocity in our dataset is $U \approx 0.7ms^{-1}$, which is almost one order of magnitude lower than in the reference studies, where $U \approx 4.5ms^{-1}$; as a result, the ratio u_τ/U remains of the same order in all cases.

Despite this order-of-magnitude consistency, the present estimate of u_τ is not sufficient to guarantee the physical reliability of the wall-unit representation. Although the snow–particle velocities appear to follow the logarithmic trend in Figure 5.4, the profile remains very oscillatory and spans a relatively narrow vertical range, which is not enough to robustly constrain the log-law slope. In addition, the snowflakes are not perfect tracer particles: their finite inertia prevents them from following the airflow as closely as the neutrally buoyant tracer particles used in laboratory water-jet experiments. The larger the particle inertia, the less accurately the motion of the snowflakes represents the underlying airflow in the boundary layer, which can cause the inferred mean horizontal velocity to deviate from the actual air velocity. This issue becomes particularly relevant in our case if the hypothesis from Section 5.1.2 is correct and only the largest and heaviest snowflakes are detected in certain regions. In that situation, the measured velocities would predominantly reflect the dynamics of a biased subset of particles rather than providing an unbiased representation of the underlying airflow.

5.2.2 Water Jet Velocity Profiles

In this section, the laboratory water-jet measurements are compared with the theoretical round-jet framework introduced in Chapter 1. This comparison is used to quantify how the differing measurement configurations in the outdoor snow and indoor water-jet experiments impact both the accuracy of the flow statistics and the reliability of the reconstructed particle tracks.

First, the axial decay of the jet centerline velocity is analysed. From Equation 1.4, we expect the centerline velocity to decrease inversely with the axial distance, so plotting $U_J/U_0(x)$ as a function of x/d should yield an approximately linear trend, similar to Figure 1.3. Figure 5.5a shows the results of the present work, plotting the centerline velocity extracted from the time-averaged Eulerian fields for all three Reynolds numbers.

The data points in Figure 5.5a lie approximately on a straight line over a wide range of axial distances, which is consistent with the expected linear increase of $U_J/U_0(x)$ in the self-similar region. At larger x/d, however, some deviations appear, most likely because

the time averaging is based on image sequences of only about 1–2 seconds, which is not sufficient to fully converge the mean centerline velocity.

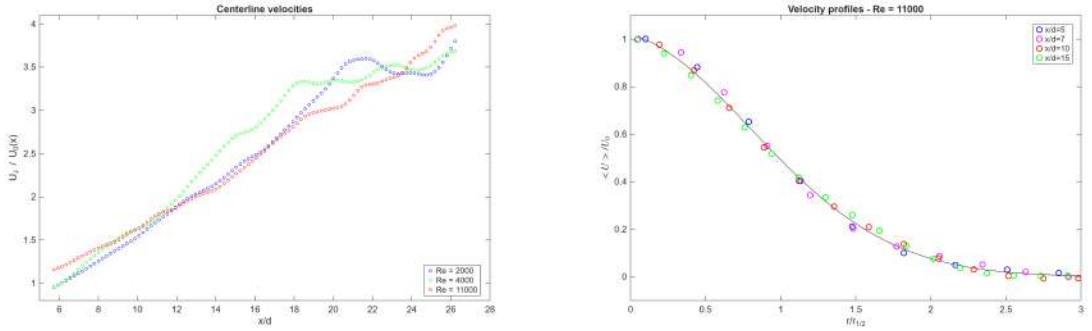

(a) (b)

Figure 5.5: (a) Axial decay of the non-dimensional centerline velocity $U_J/U_0(x)$ as a function of the streamwise distance x/d for the three Reynolds numbers investigated. (b) Normalised radial profiles of the mean axial velocity $\langle U \rangle / U_0$ as a function of the scaled radius $r/r_{1/2}$ for Re = 11000, compared with the theoretical similarity profile (black line).

In addition to the centerline decay, we also analyse the radial structure of the mean axial velocity, which was introduced in Chapter 1 as a key signature of the self-similar round-jet solution. Figure 5.5b illustrates the radial profiles of the mean axial velocity for the highest Reynolds number case, $Re = 11000$, at four different downstream locations; the velocity is scaled by the local centerline value $U_0(x)$, and the radial coordinate is scaled by the corresponding half-width $r_{1/2}(x)$ to obtain a non-dimensional representation directly comparable to the similarity profiles shown in Figure 1.4c. The resulting plot behaves as expected: all data points from the normalised centerline profiles collapse onto a single curve that closely follows the Gaussian-like similarity profile, shown as the black line in the figure.

These results confirm that the laboratory water jet behaves as a canonical turbulent round jet: both the centerline decay and the self-similar radial profiles are in very good agreement with the theoretical model. This provides a robust benchmark case in which the reconstructed tracks and Eulerian fields clearly satisfy the expected flow physics, in contrast to the outdoor snow experiment, where the more challenging measurement configuration leads to larger uncertainties and less reliable flow statistics.

Chapter 6

Conclusion

This final chapter concludes the present work by summarising the process and findings of both the snow-tracking and water-jet experiments and by stating the conclusions that emerged from the discussions in the previous chapter. Following the comparative structure adopted throughout the thesis, the two configurations are first summarised together and then compared, before assessing the potential feasibility of an STB-based snow tracking method and outlining how this work could be continued in future projects.

A snow-tracking experiment was carried out in Davos over a $10 \times 10 \times 10\ m^3$ measurement volume imaged by sixteen cameras arranged in four arrays. A drone-based calibration, relying on the pinhole projection model of Equation 3.3, established the mapping between object and image spaces. A tracking algorithm developed at IFD was applied to the images to reconstruct around 700'000 Lagrangian snowflake trajectories over four time steps, providing the results presented in Chapter 4. In a second step, a water-jet experiment was performed at the DLR laboratory in Göttingen, where an upward water jet was seeded with 50 μm particles and imaged by four cameras at three Reynolds numbers. The OTF and camera calibration were performed with the DaVis software from LaVision, using a target-based approach with a third-order polynomial mapping, and were further refined through a volume self-calibration. The Lagrangian particle trajectories were then reconstructed by means of the STB algorithm described in Section 1.4, providing the results presented in Chapter 4.

The accuracy and physical coherence of these results were assessed from several points of view. For both experiments, the reprojection error was first analysed, which needs to be minimised, especially when the STB algorithm is used, due to its high dependence on residual images. For the water-jet experiment, an extremely accurate calibration was achieved, reaching errors on the order of a few hundredths of a pixel, which corresponds to an almost perfect reconstruction. On the other hand, the mean errors recovered from the projection-matrix-based reprojection are significantly higher, varying between four

and five pixels depending on the camera, as reported in Table 4.3. A Soloff mapping was then applied to the dataset in an attempt to reduce these errors, yielding successful results: fifth- or sixth-order polynomial functions decreased the reprojection error to between two and three pixels, as reported in Table 4.4. Even though sub-pixel accuracy was not achieved, a single snowflake spans several pixels in the camera images, so the residual error remains smaller than the typical snowflake radius, resulting in very promising sub-snowflake accuracy.

The physical consistency of the reconstructed tracks was then discussed. STB has already been applied to track tracer particles in water jets on several occasions [20, 24], proving the feasibility of the task and the reliability of the resulting Lagrangian trajectories. It is, therefore, not surprising that the tracks recovered in the present work, shown in Figure 4.12, appear to represent the water jet correctly. However, as discussed in the previous chapter, the high particle density cases exhibit, in the last centimetres of the volume, some high-speed particles that are not physically feasible since no walls or perturbations could induce such a motion; these are, therefore, ghost tracks. The advantage in this case is that they are easy to identify and can be removed by applying a simple velocity threshold. Looking at the reconstructed snow tracks shown in Figure 4.6, the particles appear to move with similar velocities, with no non-physical high-speed tracks. A tempting conclusion could be that the algorithm did not reconstruct any ghost tracks and that all the detected particles are real; unfortunately, this is not the case. As shown in Section 5.1.3, particles can be reconstructed incorrectly in such a way that the resulting ghost track exhibits a velocity similar to that of the surrounding flow, which is especially true in coherent flows. These ghost particles are therefore particularly difficult to identify. The average number of occlusions reported in Table 4.3 confirms that many of these line-of-sight particles were reconstructed within the measurement volume, introducing a problematic bias in the statistics. A possible solution to this problem would be to evaluate the tracks over more than four frames and to include the fourth view from the first camera array, which was not used for the reconstruction due to initialisation problems.

As a final point, the accuracy and correctness of the results were assessed by reconstructing some flow physics in order to verify how the reconstructed tracks and Eulerian fields agree with the theory introduced in Chapter 1. For the water jet, the axial decay of the centerline velocity was plotted, together with several normalised velocity profiles shown in Figure 5.5, all of which follow the expected behaviour almost exactly. This further confirms that the STB algorithm, combined with an accurate measurement and calibration setup, yields a physically correct reconstruction, as expected, even at high particle densities. The same procedure was applied to the snow experiment results, where the curves describing the three regions of the ABL were plotted. The logarithmic law of Equation 1.2, which describes the behaviour in the logarithmic layer where the measurement volume was located, was fitted to the mean horizontal velocity profile recovered

from the experiment, as shown in Figure 5.4. The fit yielded an estimate of the friction velocity that differs from previous works; however, as discussed in the previous chapter, this difference can be explained by a corresponding difference in the average horizontal flow speed.

6.1 Snow Tracking STB Feasibility

A main goal of this thesis was to assess whether the STB algorithm, originally developed for controlled laboratory experiments at high particle image densities, could be applied to the much more challenging case of natural snowfall tracking. The results discussed in the previous chapters provide several elements to evaluate this feasibility.

The initial concern was the difficulty of obtaining an accurate calibration due to the large measurement volume and the complex outdoor conditions, which would result in inaccurate reprojections and make the implementation of a STB algorithm on such a complex dataset very unlikely. However, the reprojection errors turned out to be sub-snowflake accurate, assuming a snow particle size of approximately 5 mm. It was shown that these errors could be further reduced by applying a Soloff mapping; the errors remain higher than those obtained in the precise water-jet experiment, but considering the particle size, sub-particle accuracy is achieved in both cases, which gives good hope for the potential feasibility of an STB implementation. What now remains to be done, and which constitutes the natural continuation of this work, is to reconstruct the optical transfer function on the snow dataset using a low particle density image and subsequently attempt the first STB processing of the snow data. Achieving this would not only enable a direct quantitative comparison with the laboratory pipeline, but would also represent, to the best of our knowledge, the first application of STB to natural snowfall.

Bibliography

[1] Robert S. Arthur and Wayne M. Angevine. "What's Next: Boundary Layer Prediction Methods". In: *Conceptual Boundary Layer: The Air Near Here*. LLNL-BOOK-804362. Lawrence Livermore National Laboratory, 2020. Chap. 5.

[2] Rafael Bölsterli. "Large-scale Imaging of Natural Snowfalls in the Field". MA thesis. ETH Zurich, 2024.

[3] Koen Muller. "Using experimental fluid dynamics to track natural snowfall". In: *Nature Reviews Earth and Environment* (Feb. 2026). ISSN: 2662138X. DOI: 10.1038/s43017-025-00748-w.

[4] Xuhui Lee. *Fundamentals of Boundary-Layer Meteorology*. 2nd ed. Springer Atmospheric Sciences. Cham: Springer, 2023. ISBN: 978-3-031-32667-7. DOI: 10.1007/978-3-031-32668-4.

[5] Cristina Pérez-Guillén, Frank Techel, Martin Hendrick, Michele Volpi, Alec van Herwijnen, Tasko Olevski, Guillaume Obozinski, Fernando Pérez-Cruz, and Jürg Schweizer. "Data-driven automated predictions of the avalanche danger level for dry-snow conditions in Switzerland". In: *Natural Hazards and Earth System Sciences* 22.6 (2022), pp. 2031–2056. DOI: 10.5194/nhess-22-2031-2022. URL: https://nhess.copernicus.org/articles/22/2031/2022/.

[6] Andras Nemes, Teja Dasari, Jiarong Hong, Michele Guala, and Filippo Coletti. "Snowflakes in the atmospheric surface layer: observation of particle–turbulence dynamics". In: *Journal of Fluid Mechanics* 814 (2017), pp. 592–613. DOI: 10.1017/jfm.2017.13.

[7] Cheng Li, Kaeul Lim, Tim Berk, Aliza Abraham, Michael Heisel, Michele Guala, Filippo Coletti, and Jiarong Hong. "Settling and clustering of snow particles in atmospheric turbulence". In: *Journal of Fluid Mechanics* 912 (2021), A49. DOI: 10.1017/jfm.2020.1110.

[8] Nathaniel Bristow, Jiaqi Li, Peter Hartford, Michele Guala, and Jiarong Hong. "Imaging-based 3D particle tracking system for field characterization of particle dynamics in atmospheric flows". In: *Experiments in Fluids* 64 (2023), p. 78. DOI: 10.1007/s00348-023-03623-4.

[9] J. R. Garratt. "The atmospheric boundary layer". In: *Earth-Science Reviews* 37.1-4 (1994), pp. 89–134.

[10] Michael Heisel. "Organization and scaling of coherent structures in the outer region of high-Reynolds-number turbulent boundary layers". PhD thesis. University of Minnesota, 2020.

[11] Stephen B. Pope. *Turbulent Flows*. Cambridge University Press, 2000. ISBN: 978-0521598866.

[12] E. R. VAN DRIEST. "On Turbulent Flow Near a Wall". In: *Journal of the Aeronautical Sciences* 23 (11 Nov. 1956), pp. 1007–1011. DOI: 10.2514/8.3713.

[13] Research Starters. *What is the Atmospheric Boundary Layer?* https://discover.hubpages.com/education/What-is-the-Atmospheric-Boundary-Layer (accessed 22 May 2026). 2021.

[14] Leonard Bauersfeld, Koen Muller, Dominic Ziegler, Filippo Coletti, and Davide Scaramuzza. "Robotics meets Fluid Dynamics: A Characterization of the Induced Airflow below a Quadrotor as a Turbulent Jet". In: *IEEE Robotics and Automation Letters* (2024).

[15] Gregory C. Burton. "The Nonlinear LES (nLES) Method: A fundamental paradigm shift in turbulence modeling". In: *Abstract*. LLNL-ABS-401873. Livermore, CA, 2008.

[16] Nicholas T. Ouellette, Haitao Xu, and Eberhard Bodenschatz. "A quantitative study of three-dimensional Lagrangian particle tracking algorithms". In: *Experiments in Fluids* 40 (2 Feb. 2006), pp. 301–313. ISSN: 07234864. DOI: 10.1007/s00348-005-0068-7.

[17] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge University Press, 2004. ISBN: 0521540518.

[18] G. E. Elsinga, F. Scarano, B. Wieneke, and B. W. van Oudheusden. "Tomographic particle image velocimetry". In: *Experiments in Fluids* 41.6 (2006), pp. 933–947. DOI: 10.1007/s00348-006-0212-z.

[19] N. A. Malik, Th. Dracos, and D. A. Papantoniou. "Particle tracking velocimetry in three-dimensional flows. Part II: Particle tracking". In: *Experiments in Fluids* 15 (1993), pp. 279–294.

[20] Daniel Schanz, Sebastian Gesemann, and Andreas Schröder. "Shake-The-Box: Lagrangian particle tracking at high particle image densities". In: *Experiments in Fluids* 57 (5 May 2016). ISSN: 07234864. DOI: 10.1007/s00348-016-2157-1.

[21] Andreas Schröder and Daniel Schanz. "3D Lagrangian Particle Tracking in Fluid Mechanics". In: *Annual Review of Fluid Mechanics* 55 (2023), pp. 511–540. DOI: 10.1146/annurev-fluid-031822-041721.

[22] Bernhard Wieneke. "Iterative reconstruction of volumetric particle distribution". In: *Measurement Science and Technology* 24 (2 2013). ISSN: 13616501. DOI: 10.1088/0957-0233/24/2/024008.

[23] Tobias Jahn, Daniel Schanz, and Andreas Schröder. "Advanced iterative particle reconstruction for Lagrangian particle tracking". In: *Experiments in Fluids* 62 (8 Aug. 2021). ISSN: 14321114. DOI: 10.1007/s00348-021-03276-7.

[24] Daniel Schanz, Massimiliano Novara, and Andreas Schröder. "Shake-The-Box particle tracking with variable time-steps in flows with high velocity range (VT-STB)". in: *Proceedings of the 14th International Symposium on Particle Image Velocimetry*. Chicago, USA, 2021.

[25] Philipp Godbersen, Sebastian Gesemann, Daniel Schanz, and Andreas Schröder. "FlowFit3: Efficient data assimilation of LPT measurements". In: *Proceedings of the*

21st International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics. Lisbon, Portugal, 2024.

[26] Sebastian Gesemann. "From Particle Tracks to Velocity and Acceleration Fields Using B-Splines and Penalties". In: (Oct. 2015). URL: http://arxiv.org/abs/1510.09034.

[27] K. Muller, C. K. Hemelrijk, J. Westerweel, and D. S.W. Tam. "Calibration of multiple cameras for large-scale experiments using a freely moving calibration target". In: *Experiments in Fluids* 61 (1 Jan. 2020). ISSN: 14321114. DOI: 10.1007/s00348-019-2833-z.

[28] Andrew W. Fitzgibbon. "Simultaneous linear estimation of multiple view geometry and lens distortion". In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2001.

[29] Stephen P. Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge: Cambridge University Press, 2004.

[30] S. M. Soloff, R. J. Adrian, and Z.-C. Liu. "Distortion compensation for generalized stereoscopic particle image velocimetry". In: *Measurement Science and Technology* 8 (1997), pp. 1441–1454.

[31] Koen Muller. "Tracking Schooling Fish in Three Dimensions: Experiments at the Rotterdam Zoo". PhD thesis. Delft, The Netherlands: Delft University of Technology, 2023. ISBN: 978-94-6419-878-2.

[32] Doruk Alp Emür. "Wavy Flow Patterns of Planar Snowfall Tracking Experiment and Scaling of Field Illumination". Mechanical Engineering Master Semester Project, Institute of Fluid Dynamics, ETH Zurich. 2025.

[33] Lucas Pavlov, Sebastien Cazin, Patricia Ern, and Veronique Roig. "Exploration by Shake-the-Box technique of the 3D perturbation induced by a bubble rising in a thin-gap cell". In: *Experiments in Fluids* 62 (2021), p. 22. DOI: 10.1007/s00348-020-03117-z.

[34] D. Schanz, S. Gesemann, A. Schröder, B. Wieneke, and M. Novara. "Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction". In: *Measurement Science and Technology* 24.2 (2013), p. 024009. DOI: 10.1088/0957-0233/24/2/024009.

[35] LaVision GmbH. *Shake-the-Box: High-Density 3D Lagrangian Particle Tracking*. Product brochure. 2024. URL: https://www.lavision.com.

[36] M. P. Langleben. "The terminal velocity of snowflakes". In: *Quarterly Journal of the Royal Meteorological Society* 80.346 (1954), pp. 174–181.

List of Figures

1.1 (a) Avalanche risk in Switzerland on 6 March 2020. Adopted from [5]. (b) A train derailed due to an avalanche on 16 February 2026 in Goppenstein, Switzerland. 1
1.2 (a) Schematic of the vertical wind speed profile in the ABL, showing the increase from the surface within the turbulent boundary layer and the quasi-uniform flow above. Reproduced from [13]. (b) Dimensionless mean velocity u^+ as a function of the dimensionless wall distance y^+. The circles, dashed line, and dot–dashed line represent experimental data at different Reynolds numbers, while the solid line shows von Driest's law. Reproduced from [11] 3
1.3 Variation of the scaled centerline velocity $U_J/U_0(x)$ as a function of the nondimensional distance x/d. Reproduced from [11]. 4
1.4 (a) Radial profiles of mean axial velocity, the dashed lines indicate the half-width. Reproduced from [11]. (b) Qualitative representation of the radial profiles of the mean axial velocity $\langle U \rangle$ of a turbulent jet flow. Adopted from [15]. (c) Scaled versions of the radial profiles of mean axial velocity; different symbols represent different axial distances. Reproduced from [11] 5
1.5 Visual interpretation of STB. The initialisation phase (with IPR-based particle detection) is highlighted in yellow, while the convergence and converged phases are shown in blue. Reproduced from [21]. 8

2.1 Overview of the experimental site in Davos, Switzerland. The tents are deployed in a 30-meter diameter arc around the light panels. 10
2.2 Main components of the volumetric imaging system. (a) Camera setup: 16 cameras were used in total, divided into 4 arrays of 4 cameras, located in the blue tents. (b) Control computer used for synchronising the cameras and recording the volumetric snowfall data. (c) Tent setup: 7 tents in total were deployed, protecting the cameras and computers during data acquisition. (d) LED light panels providing the illumination necessary to detect snowflakes. 11

2.3 (a) Large scale snow tracking experiment in low wind condition. (b) Calibration flight. The four light bulbs mounted on the drone follow a predetermined path through the measurement volume and are recorded by the cameras for calibration purposes. (c) Overview of the illuminated measurement volume and tent arrangement during high wind conditions. 13
2.4 (a) General setup of the water jet experiment. Four cameras are positioned around the hexadecagonal-sided tank, pointing at the light beam to capture illuminated tracer particles emerging from the jet nozzle at the bottom. (b) Close-up view of the illuminated measurement volume inside the tank, showing the light beam generated by the optical system. 14
2.5 Optical components. (a) Converging lens positioned above the water surface, with a glass panel underneath to eliminate refraction caused by surface waves generated by the flow. (b) Aperture (passepartout) mounted on top of the lens to block stray light outside the optical path. (c) Optical setup for generating the light beam. Light from two LEDs is combined using a prism to create a uniform illumination volume for particle tracking. 15
2.6 (a) Close-up view of the four-camera array positioned around the measurement volume. (b) Freely moving calibration target with markers distributed on two planes to provide a shallow depth. 16

3.1 (a) Drone trajectory used for the calibration flight, reconstructed in MATLAB from the three-dimensional coordinates over the measurement volume. (b) Drone velocity along the path. 19
3.2 Camera calibration in DaVis using a freely moved target. (a) Target image recorded by one camera (Camera 4), with the fitted calibration grid overlaid on the detected marker points. (b) Corresponding calibration mapping parameters for plane 6 (centre of the measurement volume). 21
3.3 VSC in DaVis. (a) Disparity map obtained from particle images on plane 6. (b) Corresponding disparity field used to update the camera-volume mapping. (c) Updated mapping parameters after VSC. 22
3.4 OTF reconstruction in DaVis using low particle density images. (a) Low-density particle image recorded by one camera. (b) Corresponding OTF particle pattern for plane 6. 23
3.5 Image preprocessing in DaVis. (a) Raw particle image before minimum subtraction. (b) Minimum intensity image subtracted from the recorded sequence. Note that the colour scales differ between the two panels. 24
3.6 Comparison of particle image density in preprocessed images. (a) Higher-density case with an average particle image density of 0.12 ppp (b) Lower-density case with an average particle image density of 0.08 ppp. 24

4.1 (a) Image of the DJI Matrice 300 light-marker tetrahedron recorded by one camera during the calibration flight. (b) Histogram of the object skew error after optimisation of the drone-based calibration. 27
4.2 (a) Histogram of the image reprojection error over all cameras and calibration frames, expressed in pixels. (b) Example distributions of the reprojection error for individual cameras within one array. 27
4.3 (a) Histogram of the sixth-order Soloff's reprojection error for camera 2. (b) Corresponding sixth-order reprojection error histograms for all cameras in array 1. 29
4.4 (a) Measured (blue) and camera-matrix-based reprojected (red) snowflake centroids on Camera 8 for frame 4. (b) Distribution of pixel-based reprojection error for Camera 8 in the same frame. 30
4.5 (a) Measured (blue) and Soloff-based reprojected (red) snowflake centroids on Camera 8 for frame 4. (b) Distribution of pixel-based Soloff's reprojection error for Camera 8 in the same frame. 31
4.6 Reconstructed snowflakes coloured by velocity magnitude in the full measurement volume (a) and in a 1 m thick central XZ-slab (b), illustrating the spatial distribution of particles inside the cube. 33
4.7 (a) Horizontal slab at the bottom of the measurement volume showing the reconstructed snowflake positions. (b) Same slab, where the particles are coloured by the Delaunay-based nearest-neighbour distance D. 34
4.8 Mean particle speed in 1 m thick slabs as a function of slab position along (a) X, (b) Y, and (c) Z for the snow tracking dataset. 34
4.9 Probability density functions (PDFs) of the three velocity components (a) U, (b) V and (c) W for the snow tracking dataset. The solid lines indicate Gaussian fits obtained from the eleven central bins of each distribution. . . 35
4.10 Time-averaged Eulerian velocity field reconstructed from the snowflake trajectories. (a) Overview of the entire measurement volume, showing velocity vectors at each grid point coloured by their magnitude. (b) One-meter-thick slab extracted from the centre of the volume along the Y-direction. 35
4.11 (a) Three-dimensional fluctuation field within the measurement volume, with vectors coloured by their instantaneous magnitude. Note that the colour scale differs from the one used for the mean velocity fields in Figure 4.10. (b) Histogram of the fluctuation magnitude for frame 4, illustrating the distribution of turbulent velocities sampled by the snowflakes. 36

4.12 Lagrangian particle tracks obtained with STB. Panels (a), (b), and (c) show the full measurement volume for Reynolds numbers $Re = 2000$, 4000, and 11000, respectively, using a track length of 15 frames. Panels (d), (e), and (f) show the corresponding central plane sections, defined by $-5 \leq z \leq 5$ mm, using a track length of 29 frames. 38
4.13 (a) Preprocessed particle image for $Re = 11000$. (b) Residual image of the same frame after STB processing. 39
4.14 STB statistics for $Re = 4000$. The number of detected particles (red), total tracks (green), and new tracks per time step (blue) are shown as a function of time step index. Panel (a) corresponds to a single-pass reconstruction, panel (b) to the multi-pass reconstruction. 40
4.15 Eulerian velocity fields reconstructed with the FineScaleReconstruction method. Panels (a), (b) and (c) show instantaneous fields for $Re = 2000$, $Re = 4000$ and $Re = 11000$, respectively, while panel (d) shows the time-averaged field for $Re = 2000$ over 2000 frames. 41

5.1 Mean normalised particle speed (blue) and number of snowflakes (orange) in 1 m thick slabs as a function of slab position along (a) X, (b) Y, and (c) Z for the snow tracking dataset. 44
5.2 (a) Schematic representation of the line of sight geometry, reported from [11]. (b) Example of an incorrectly reconstructed particle due to line of sight ambiguities. 45
5.3 Probability density functions of the particle speed in the water-jet experiment, evaluated from particles located in the last 20 mm of the measurement volume for (a) $Re = 2000$ and (b) $Re = 11000$; the insets show the same distributions on a semi-logarithmic scale. 47
5.4 Law-of-the-wall representation of the mean horizontal snow–particle velocity profile in dimensionless viscous units. (a) Linear scale, showing the measured data (red) together with the viscous sublayer, logarithmic law, and von Driest interpolation. (b) Logarithmic representation of the same profile . . 48
5.5 (a) Axial decay of the non-dimensional centerline velocity $U_J/U_0(x)$ as a function of the streamwise distance x/d for the three Reynolds numbers investigated. (b) Normalised radial profiles of the mean axial velocity $\langle U \rangle / U_0$ as a function of the scaled radius $r/r_{1/2}$ for Re = 11000, compared with the theoretical similarity profile (black line). 50

List of Tables

2.1 Overview of the snowfall imaging datasets measured during the 2024/2025 and 2025/2026 campaigns. For each date, the type and number of datasets, as well as which drone has been used for the calibration, are reported. . . . 14

4.1 Mean and standard deviation of the Soloff reprojection error for camera 2 as a function of the polynomial order. The statistics are computed from 1044 calibration samples. 28
4.2 Mean and standard deviation of the Soloff reprojection error for each camera, together with the number of calibration samples used for each fit and the mean base reprojection error. Values highlighted in red indicate poor performance of the Soloff mapping for some cameras. 29
4.3 Number of snowflake samples, pixel-based reprojection statistics (mean and standard deviation) and average number of occlusions for each camera in frame 4. 31
4.4 Polynomial order, number of snowflake samples, and pixel-based reprojection statistics (mean and standard deviation) obtained with Soloff's mapping for each camera in the selected frame. 32
4.5 Number of reconstructed snowflakes per 1 m thick slab along the three coordinate directions. 33

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every written paper or thesis authored during the course of studies. **In consultation with the supervisor**, one of the following two options must be selected:

☐ I hereby declare that I authored the work in question independently, i.e. that no one helped me to author it. Suggestions from the supervisor regarding language and content are excepted. I used no generative artificial intelligence technologies[1].

☒ I hereby declare that I authored the work in question independently. In doing so I only used the authorised aids, which included suggestions from the supervisor regarding language and content and generative artificial intelligence technologies. The use of the latter and the respective source declarations proceeded in consultation with the supervisor.

Title of paper or thesis:

Intercomparison of Snow Imaging and Water Jet Tracking Experiments

Authored by:
If the work was compiled in a group, the names of all authors are required.

Last name(s):	**First name(s):**
Cavin	Jules

With my signature I confirm the following:
- I have adhered to the rules set out in the Citation Guidelines.
- I have documented all methods, data and processes truthfully and fully.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for originality.

Place, date	**Signature(s)**
Wipkingen, 15.07.2026	

If the work was compiled in a group, the names of all authors are required. Through their signatures they vouch jointly for the entire content of the written work.

[1] For further information please consult the ETH Zurich websites, e.g. https://ethz.ch/en/the-eth-zurich/education/ai-in-education.html and https://library.ethz.ch/en/researching-and-publishing/scientific-writing-at-eth-zurich.html (subject to change).