Light reflections detection and correction for Robotic Volumetric PIV

Laura Porcar Galan

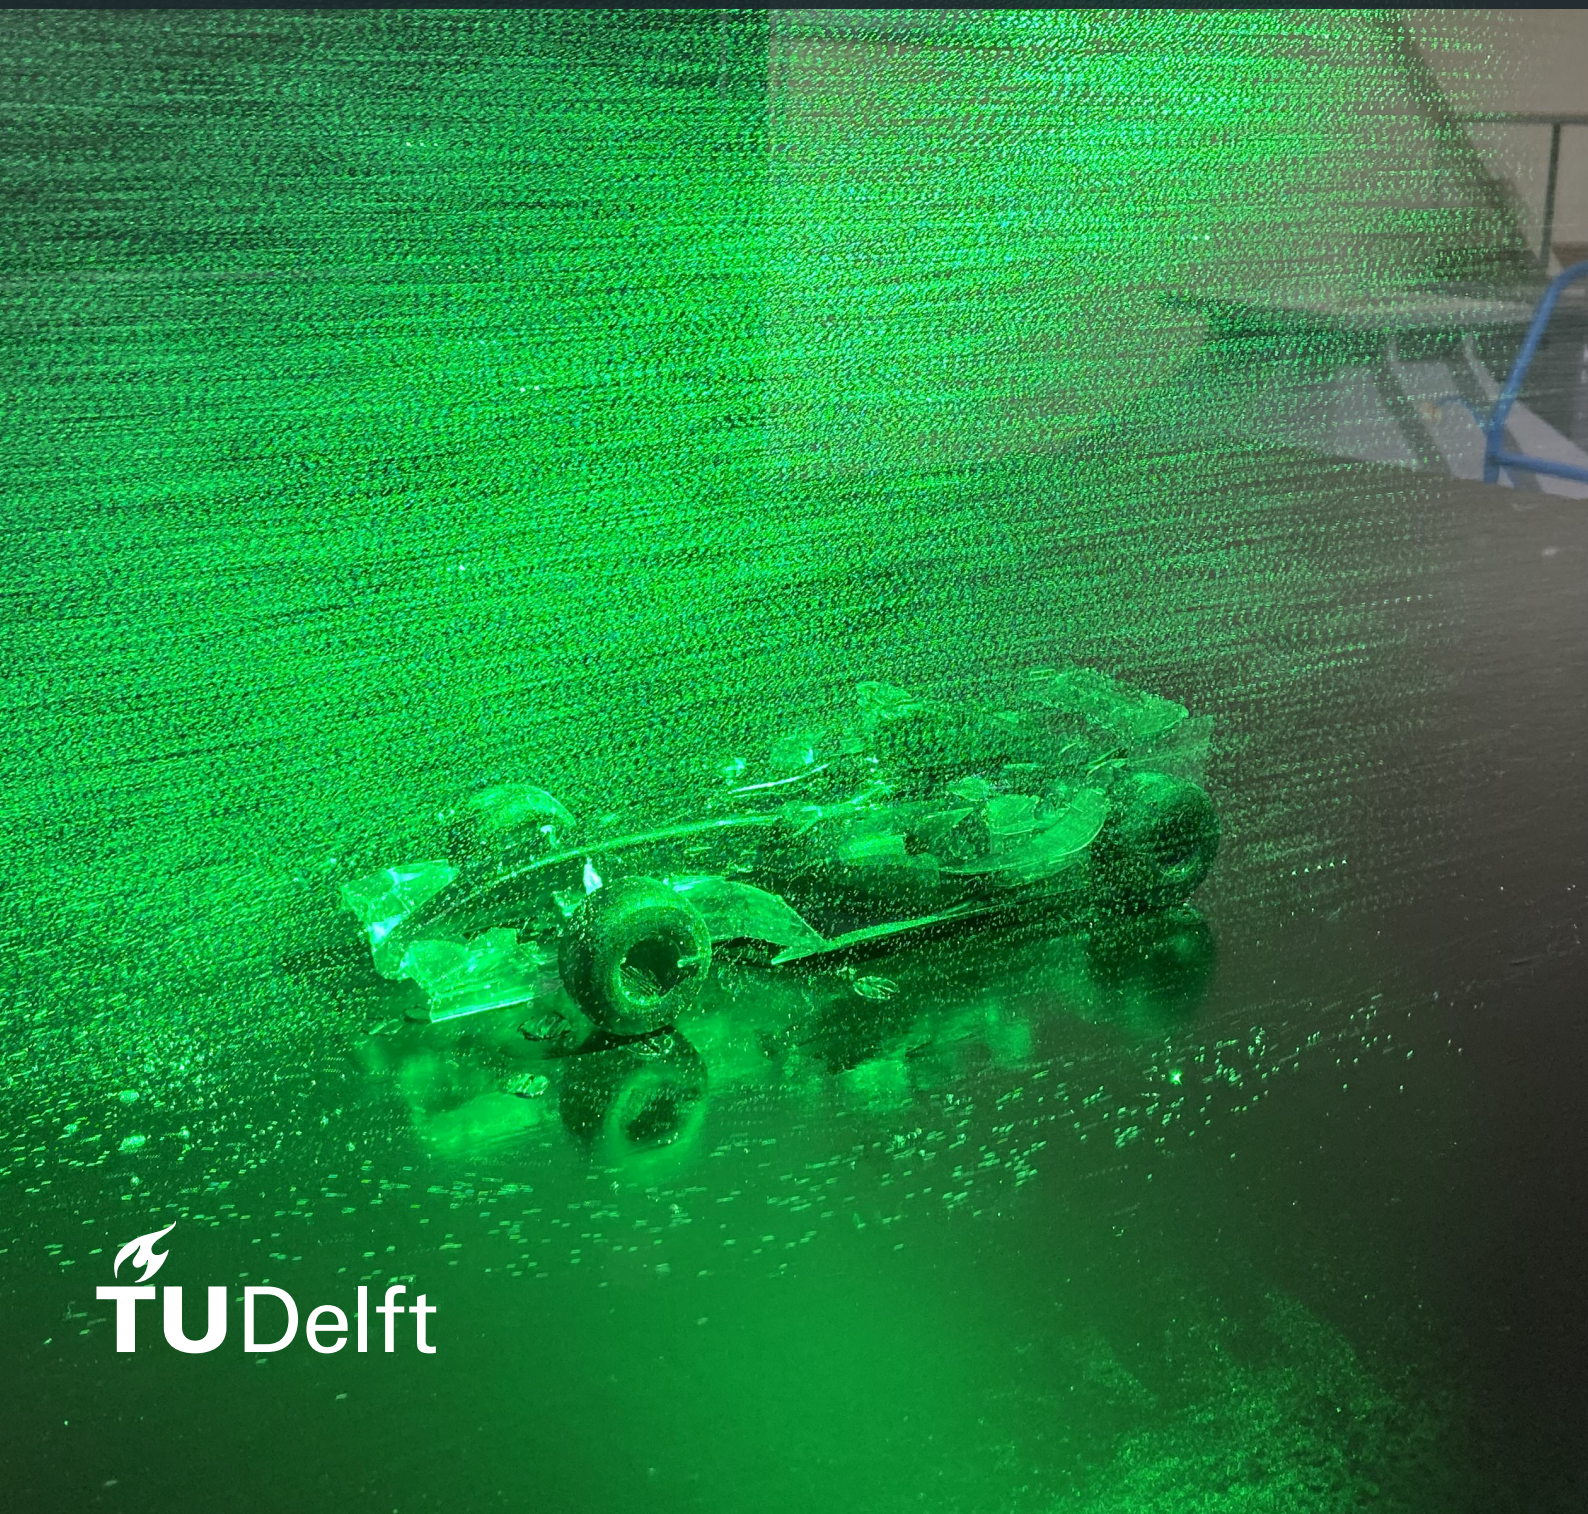

TU Delft

Light reflections detection and correction for Robotic Volumetric PIV

Thesis report

by

Laura Porcar Galan

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on June 17^{th}, 2024 at 13:30

Student number:	5609216
Project Duration:	September, 2023 - June, 2024
Place:	Faculty of Aerospace Engineering, Delft
Thesis committee:	
Chair:	Dr. Wei Yu
Supervisors:	Dr. Andrea Sciacchitano
	Prof. Dr. Fulvio Scarano
	Adrian Grille Guerra
External examiner:	Dr. Tercio Lima Pereira

An electronic version of this thesis is available at http://repository.tudelft.nl/.

Faculty of Aerospace Engineering · Delft University of Technology

Preface

The master thesis completion marks the end of my student life (for now), marking the ending of years of hard work, dedication and countless learning experiences, with up and downs every year. Completing this master's thesis has been a significant milestone in my academic journey and I am grateful for the opportunity to explore deeply a research topic of my chosen field of study.

I would like to express my sincere gratitude to my supervisors, Andrea Sciacchitano and Fulvio Scarano, for guiding me throughout the duration of this thesis. Their expertise and insights have been fundamental and extremely valuable in shaping this work and guiding me through all its stages. I must also especially thank Adrian Grille Guerra for being so attentive and helping during the whole duration of the thesis. Without his help and advice, this project would not have been the same.

I would also like to extend my deeply felt thanks to my family and friends for their constant encouragement and support. Their belief in me has been a constant source of strength and motivation and I am truly grateful for their presence in my life. *Gracias a todos*.

This thesis is as much a reflection of their support and guidance as it is of my own efforts, and I am thankful for their contributions to my journey.

Laura Porcar Galan

14^{th} May 2024

Abstract

Laser light reflection mitigation in Particle Image Velocimetry (PIV) is crucial for accurate flow field measurements. While numerous methods exist for planar PIV, fewer have been developed for volumetric PIV systems, especially for coaxial setups like Robotic PIV. Light reflections in volumetric PIV experiments result in high-intensity regions that corrupt particle detection and analysis.

This study presents three novel approaches for treating light reflections in Robotic PIV experiments. The first and second methods use image filtering and masking techniques in the wavenumber space to separate particle images from reflection regions. The first technique called *Spatial Fourier Filter* involves decomposing the image signal into low- and high-wavenumber components using the 2D discrete Fourier transform (DFT). A high-pass filter is then applied to attenuate the intensity of reflection regions. Then, the second methodology *Spatial Fourier Filter + Mask* takes the resulting image from the first method and performs a step of automated adaptive masking to remove residual reflection areas that the filtering approach is not able to eliminate. The third methodology named *3D-based Particle Concentration Mask* acts in a later stage of the processing pipeline, creating a 3D mask on the instantaneous processed Shake-the-Box data by analysing the particle concentration distribution over the flow domain.

The proposed methods are tested on experimental data obtained from experiments performed with Robotic PIV on three different geometries: a side-view mirror, Formula 1 car and a propeller. The tests were conducted at one of TU Delft Aerospace Engineering Faculty's facilities, the W-tunnel in the High-Speed Laboratory (HSL). Comparison between raw and pre-processed images, as well as particle tracking results, is presented.

The results from this data comparison show unsatisfactory outcomes from both Spatial Fourier Filter and 3D-based Particle Concentration Mask, which fail to fully remove the spurious regions. Nevertheless, the results confirm the successful removal of reflection-induced artifacts in instantaneous images by using the spatial Fourier filter automated masking approach. The developed image pre-processing strategy effectively removes reflection regions in Robotic PIV images, preventing the appearance of spurious particle tracks. The method shows promising results mitigating unsteady light reflections in Robotic PIV, improving the accuracy of flow field measurements. Additional attention is required in the PIV sequence creation step to ensure an adequate level of overlap between measurement volumes. This facilitates addressing the spatial gaps introduced by the masking procedure, that have been proven to robustly be filled in by the multi-view advantage offered by Robotic PIV.

Contents

Nomenclature

List of Abbreviations

1D	One Dimension
2C	Two Components
2D	Two Dimensions
3C	Three Components
3D	Three Dimensions
4D	Four Dimensions
CCD	Charge Coupled Device
CCMOS	Complementary Metal-Oxide Semiconductor
CVV	Coaxial Volumetric Velocimetry
DFT	Discrete Fourier Transform
F1	Formula 1
FFT	Fast Fourier Transform
FOV	Field Of View
fps	Frames per second
FT	Fourier Transform
HFSB	Helium-Filled Soap Bubbles
HPF	High-Pass Filter
HWA	Hot-Wire Anemometry
IPR	Iterative Particle Reconstruction
LDA	Laser-Doppler Anemometry
Nd:YAG	Neodymium-doped Yttrium Aluminium Garnet
Nd:YLF	Neodymium-doped Yttrium Lithium Fluoride
OJF	Open-Jet Facility
OTF	Optical Transfer Function
PDF	Probability Density Function
PIV	Particle Image Velocimetry
POD	Proper Orthogonal Decomposition
PSD	Power Spectral Density
PTV	Particle Tracking Velocimetry
ROM	Reduced Order Model
RPM	Revolutions Per Minute
SFF	Spatial Fourier Filter
SFFM	Spatial Fourier Filter + Mask
SNR	Signal-to-Noise Ratio
SSIM	Structural Similarity Index Measure
STB	Shake-the-Box
VAWT	Vertical Axis Wind Turbine

List of Symbols

Δt	Time step
δz	Focal depth or depth of field
$\dot{N}$	Effective bubble production rate of the seeding rake [particles/s]
$\hat{\epsilon}_G$	Sub-pixel particle peak position [px]
κ	Spatial frequency (or wavenumber) [px^{-1}]
κ_f	Filter's cut-on wavenumber [px^{-1}]
κ_x	Spatial frequency (or wavenumber) in horizontal direction
κ_y	Spatial frequency (or wavenumber) in vertical direction

λ	Wavelength [nm or px]]
$\mathcal{F}$	Fourier transform
σ	Standard deviation
φ	Phase angle [rad]
A_{rake}	Area occupied by the seeding rake [cm^2]
C_{HFSB}	HFSB particle concentration
D	Aperture
d_τ	Particle effective diameter
f	Temporal frequency [Hz] (or focal length when referring to camera characteristics)
$f_\#$	f-number or f-stop
G	Fourier transform
g	Signal function in the real domain
H	Fourier transform filter
L_r	Reflection length
M	Magnification
n	Wind tunnel contraction ratio
P	Power spectrum
s	Separability criterion
t	Time [s]
u	Forward velocity component
U_∞	Free-stream velocity [m/s]
v	Sideways velocity component
w	Vertical velocity component
x	Forward component
y	Sideways component
z	Vertical component

List of Figures

List of Tables

1 Introduction

One of the main challenges in the aerospace industry is the characterization of flow fields. By understanding how air or other fluids flow around objects, engineers can optimize the design for improved performance and efficiency under different conditions. Additionally, it is essential for ensuring the safety and stability of vehicles and structures, particularly in cases of high speeds or turbulent flow conditions. Knowing the flow field characteristics allows to predict aerodynamic forces, like the drag and lift forces, from which engineers can improve fuel consumption efficiency (drag-reduction techniques), stability and overall performance. Furthermore, acquiring experimental data is essential for validating and improving computational fluid dynamics (CFD) models, allowing to improve their accuracy and reliability. Therefore, understanding the properties of a flow field is essential for optimizing designs, predicting system behaviour and improving the efficiency and reliability of systems in various engineering applications.

Among a wide variety of methodologies for flow field characterization, Particle Image Velocimetry (PIV) has gained more importance in the past few years, becoming one of the most promising flow measurement techniques. PIV is a non-intrusive flow measurement technique that allows to retrieve quantitative information of the velocity field at a certain time and location. Its simplest set-up is the 2D or Planar PIV, which measures the velocity within a two-dimensional plane using a single camera. However, the vast majority of flows are 3D in nature, like turbulent flows.

Robotic PIV is an example of a technique that allows to retrieve the 3D volumetric velocity field. The process involves capturing 2D images from three or more camera views and then reconstruct the 3D particle field from the individual images of an illuminated measurement volume. One of the main advantages is its flexibility and adaptability in camera positioning, that allows to study complex geometries or experimental setups with difficult access. In short, the pipeline of Robotic PIV includes first a system calibration, then when everything is set to start, the images are acquired. These are pre-processed and then post-processed with a volumetric particle tracking algorithm called Shake-the-Box, which allows to obtain the velocity field in the measurement volume.

The problem arises when the laser light hits on a solid surface, resulting in unwanted light reflections. These reflections appear with an intensity higher than that of the particle images, yielding a considerable issue during data processing. The regions where reflections appear often exhibit a complete absence of data, creating voids within the flow field information. Moreover, they can introduce inaccuracies in the obtained flow field velocity, leading to potentially erroneous interpretations of the flow dynamics. This presents a challenge, particularly due to the loss of information caused by the incidence angle of the camera perspective capturing these reflections. The impact of reflections becomes even more detrimental when they do not present themselves as

data gaps but rather as misleading velocity values within the flow field, significantly compromising the accuracy and reliability of the interpretation.

Consequently, it is essential to develop a robust methodology capable of detecting, not only the presence of spurious regions in the 2D images captured in the camera plane, but also tackle and control their impact on the epipolar lines within the 3D reconstructed volume. Investigating and developing such methodology is crucial to address the challenges posed by reflections in the Tomographic or Robotic PIV process, aiming to improve the accuracy and reliability of flow field analysis and interpretation.

The main objective of this research project is to *minimize the erroneous data due to reflections and ensure an acceptable particle tracking by defining a procedure to detect and mitigate reflections in Robotic PIV data*.

1.1. Research Formulation

The main research question of this project would be:

- **(RQ.) How can a methodology effectively identify and mitigate the effects of spurious regions of reflections in Robotic PIV across various PIV datasets?**

From the latter, a set of sub-questions can be considered:

- (RSQ1) Which image pre-processing technique can be applied to detect and mitigate reflections by only using an instantaneous image?
- (RSQ2) How do reflections appear in the 3D velocity field? Is there a distinguishable characteristic that could be attributed to and, hence used to remove this region of the data?
- (RSQ3) What are the advantages of the developed methodologies against state-of-the-art techniques in identifying and mitigating spurious regions?
- (RSQ4) What are the potential challenges and limitations in implementing the proposed methodology and how can these be addressed?

1.2. Report structure

The structure of the report is as follows. First, Chapter 2 presents a literature review, discussing the fundamentals of PIV, a brief summary of the state-of-the-art reflection treatment techniques and information about the implementation of Fourier analysis in image pre-processing. The report then delves into Chapter 3 where a characterization of light reflections is exposed before presenting the three proposed methods for reflection mitigation in Chapters 4.1, 4.2 and 4.3, respectively, where each method is detailed, including its theoretical background and implementation steps. An experimental campaign description follows in Chapter 5, outlining the setup and procedure used to validate the proposed methods. The results of the experimental campaign are then presented in Chapter 6, evaluating the effectiveness of each method and discussing their impact on flow field analysis accuracy. Finally, the report concludes with Chapter 7 with a summary of the findings and their implications for Robotic PIV experiments, along with recommendations for future research in this area.

2
Literature Review

2.1. Particle Image Velocimetry

Fluid visualization by introducing smoke or paint in the test section is usually used to observe the behaviour of the flow in a qualitative way. Moreover, there are other techniques like Hot-Wire Anemometry (HWA) or Laser-Doppler Anemometry (LDA) that provide the velocity information at a specific point. However, these are unable to capture the instantaneous behaviour of the flow field, as would be useful in the case of turbulent flows. Consequently, Particle Tracking Velocimetry (PTV) and Particle Image Velocimetry (PIV) appeared to fill this need. In PTV, individual particles are tracked (Lagrangian framework), hence it is restricted to low-seeding density images. Conversely, PIV is a flow measurement technique that aims to provide the instantaneous velocity field in a cross-section of the flow, determining the velocity of the fluid within finite rectangular areas (Eulerian framework). Therefore, the latter technique can be applied to higher density images, allowing to obtain the velocity field of larger areas. In Figure 2.1, the difference in the spatial resolutions between both methodologies is shown.

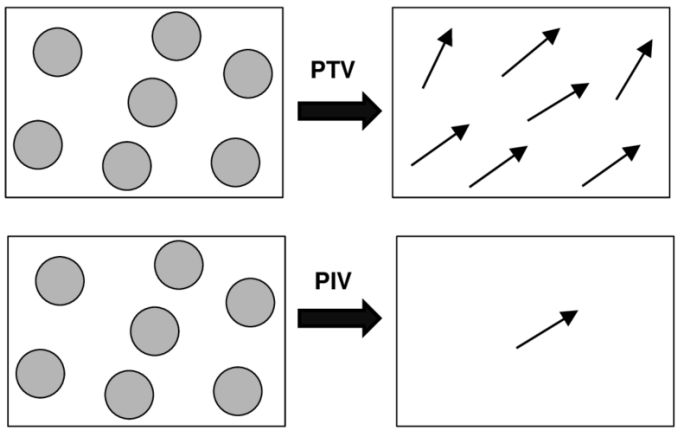

Figure 2.1: Comparison of spatial resolution between PTV and PIV. Source: [1].

Both techniques rely on the definition of velocity:

$$u = lim_{\Delta t \to 0} \frac{\Delta s}{\Delta t} \tag{2.1}$$

where Δs is the displacement and Δt the time step separation between two consecutive images acquired. PTV is restricted to low-speed flows, for instance in cases of microscale flows or flows

with separation or recirculation regions. PIV is mostly used in the research field, and due to the continuous development and innovation, its popularity on the application in the industrial environment is rising. For instance in the development of engines, automotive aerodynamics, wind turbines, sprays, multi-phase flows, etc [2].

Figure 2.2 shows an example of set-up for planar PIV able to measure two velocity components within the flow field in a plane in a wind tunnel (2C-2D PIV) [3].

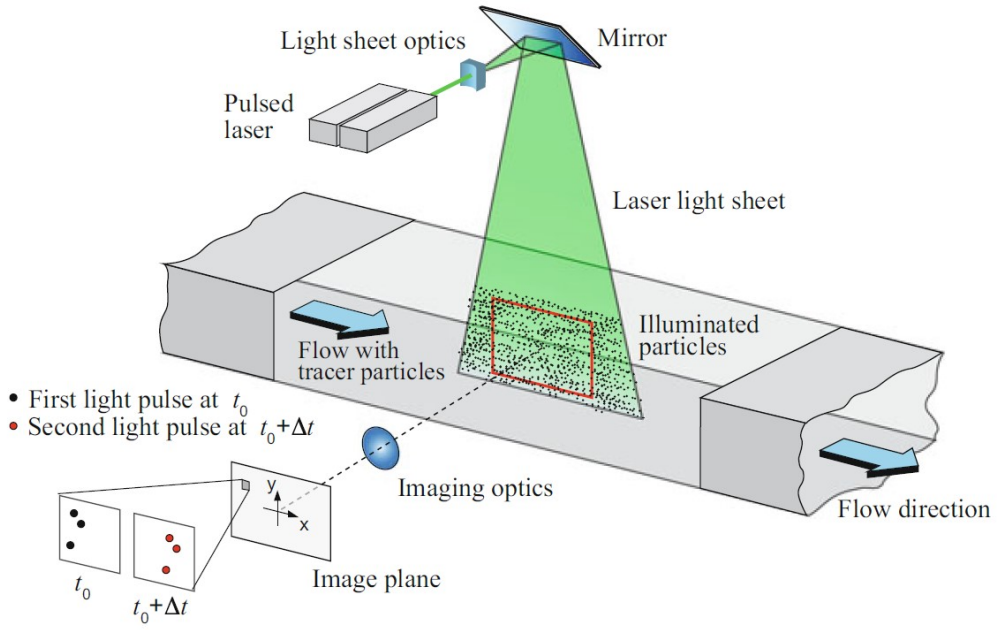

Figure 2.2: Skecth of a set-up for 2C-2D PIV in a wind tunnel. Source: [3].

Tracer particles are introduced into the flow domain. Table 2.1 shows the typical seeding particles used in PIV.

Table 2.1: Technical specifications of seeding particles. Source: [4]

Fluid	**Material**	**Diameter (μm)**	**Density (kg/m^3)**
Air	DEHS	$1-3$	10^3
	Glycol-water solution	$1-3$	10^3
	Vegetable oil	$1-3$	10^3
	TiO_2	$0.2-0.5$	$1-4\cdot10^3$
Water	Latex	$5-50$	10^3
	Sphericell	$10-100$	$0.95-1.05\cdot10^3$
	Silver coated hollow glass spheres	$30-100$	$>10^3$

The size of these particles is not arbitrary. When compared to air, the density of the particles is normally much higher than that of the fluid. Therefore, the particle diameter must be chosen sufficiently small so that the flow is minimally disturbed (PIV is considered a non-intrusive technique because of this) and they can properly follow the flow motion. However, there is an aspect that opposes this requirement: the particles also need to scatter enough light to be visible and, hence

captured by the recording device. The efficiency in how well the particle scatters the light depends on: the ratio of the refractive indexes $n = c/c_0$ (where c the local speed of light and c_0 is the speed of light in vacuum) of the particles to that of the fluid, the wavelength of the light source used λ and the particle diameter d_p. The larger the particle diameter, the more light will scatter, but, at the same time, the higher the danger to disturbing the flow field. Therefore, there is a clear trade-off when choosing the right particle size. An alternative option is to introduce a fluorescent dye to the tracer particles. In this case, when the dye in the particle absorbs the laser light, it will emit it at a longer wavelength, which facilitates distinguishing between particles from other objects or artefacts in the environment. Adding a dye into the particles can be suitable when used in water (products with density close to water is achievable), but in the case of air, it can be challenging to obtain similar density such that the buoyancy is preserved, hence this approach is not preferred.

In order to make the tracer particles visible and clearly detected by imaging devices, lasers are used as illumination sources. Some of the properties that make them suitable for PIV purposes are that they can produce a pulsed, collimated (parallel to each other) and monochromatic light beam that can be easily shaped into a thin light sheet. The most popular lasers commonly used with PIV are listed in Table 2.2.

Table 2.2: Technical specifications of three most popular lasers. Source: [3]

Type	Wavelength [nm]	Power or pulse energy	Repetition rate [Hz]
Argon ion	514, 488	10-30 W	CW
Nd:YAG	532	320 mJ	10
Nd:YLF	526	10-30 mJ	10 - 10000

The most common light source used for PIV experiments is the solidstate frequency-doubled neodymium-doped yttrium aluminium garnet laser (Nd:YAG). Along with this light source, lenses are added to the set-up to generate the appropriate laser sheet to illuminate the desired volume allowing keeping the sheet thickness constantly small (see Figure 2.2). The laser is fired at least twice within a very short and known time interval Δt. This Δt will then be used to obtain the displacement of the particles between the consecutive frames taken by the recording device(s).

Commonly a CCD or CMOS camera conforms the recording system. The acquisition rate (frames per second, fps in Hz) has to be sufficient enough to enable recording consecutive images given the Δt, time separation between laser pulses. Table 2.3 shows the specifications of two cameras typically used in PIV.

Table 2.3: Technical specifications of two examples of recording systems. Source: [3]

Type	Sensor	Pixel size	Frame rate	Exposure time
LaVision sCMOS	16-bit	6.5μm	0.5-0.1 kHz	1μs - 100 ms
Photron Fastcam SA1.1	12-bit ADC	20μm	1-675 kHz	1/fps to 1μs

The imaging system is characterized by its focal length f, f-number (or f-stop) $f_\# = f/D$ (where D is the lens aperture diameter) and image magnification M. The latter is defined as the ratio of d_i the distance between the image plane and lens and d_o the distance between the lens and the object plane.

$$\frac{1}{d_i} + \frac{1}{d_o} = \frac{1}{f} \qquad\qquad M = \frac{d_i}{d_o} \tag{2.2}$$

Ideally, if the tracer particles were captured by the imaging system as *spots*, the particle images would present a diameter of:

$$d_{geom} = M d_p \tag{2.3}$$

where d_p is the effective particle diameter and M is the magnification.

Nevertheless, imaging systems are affected by diffraction [5]. This effect occurs when plane light waves impinge on an opaque screen containing a circular aperture and, in consequence, a far-field diffraction pattern is generated as shown in Figure 2.3.

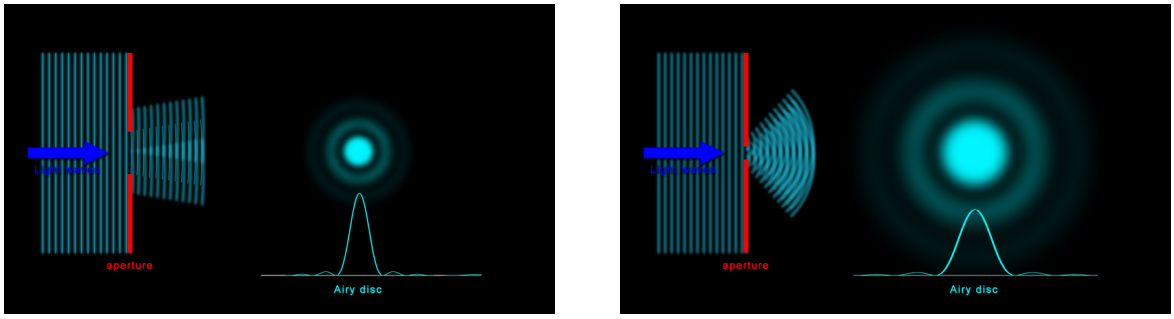

(a) **(b)**

Figure 2.3: Diffraction example with (a) larger and (b) smaller apertures. Source: [6].

The use of a lens allows the far field pattern to be imaged on the image sensor. However, when a point is to be imaged (e.g. a small scattering particle inside the light sheet), it does not appear as a point in the image plane but appears as a diffraction pattern even if it is imaged by a perfectly aberration-free lens. The pattern imaged results in a central peak of the intensity distribution called *Airy disk*, where the rings around the maximum showing decreasing brightness are called *Airy rings* (Figure 2.4). The light distribution in the Airy disk is well approximated by a

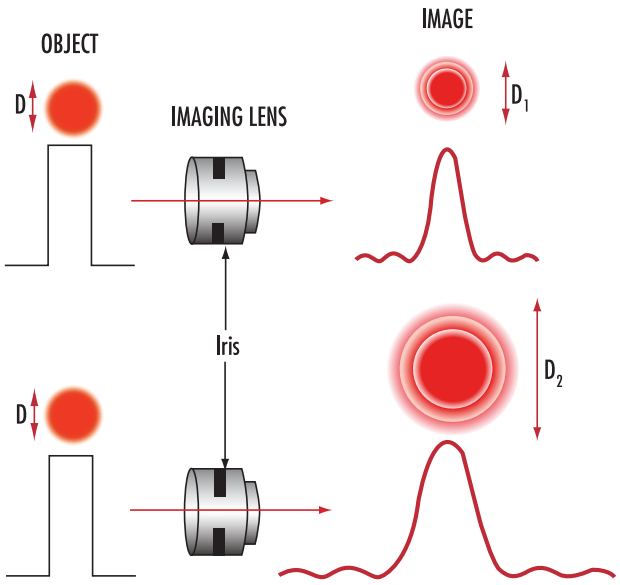

Figure 2.4: Airy disk as aperture changes. (Top) $f_{\#}$ = 2.8; (bottom) $f_{\#}$=8. Source: [7].

Gaussian intensity distribution $I(x)$ [8].

$$I(x) = I_{max} e^{\left(-\frac{(x-x_o)^2}{2\sigma^2}\right)} \tag{2.4}$$

where I_{max} is the maximum value of intensity, σ represents the radius of the Gaussian-like intensity profile and must be set to $\sigma = \lambda(1+M)f_{\#}\sqrt{2}\pi$, in order to approximate diffraction limited imaging [3].

Considering the imaging system configuration and the size of the airy disk, the diffraction effect that limits the minimum particle image diameter can be accounted with the following expression:

$$d_{diff} = 2.44\lambda(1+M)f_{\#} \tag{2.5}$$

The diameter of the Airy disk d_{diff} represents the smallest particle image that can be obtained for a given imaging system. Hence, an acceptable estimate of the particle image diameter is given by the Euclidean sum of the geometric particle diameter and the diameter due to diffraction effects:

$$d_{\tau} = \sqrt{d_{geom}^2 + d_{diff}^2} \tag{2.6}$$

In PIV, d_{diff} is more predominant when recording small particles ($\approx$ few microns) at small magnifications. On the contrary, the geometric diameter d_{geom} becomes more significant in the case of larger particles and/or larger magnifications. As depicted in Figure 2.3, the size of the Airy disk depends on the aperture. A smaller aperture will produce a larger airy disk compared to a larger aperture.

The particle image diameter is defined by Equation 2.6 only when particle images are in focus, i.e. when the light sheet thickness Δz_0 is smaller than the focal depth or depth of field δz of the optical system given by:

$$\delta z = 2f_{\#}d_{diff}\frac{M+1}{M^2} = 4.88f_{\#}^2\left(\frac{M+1}{M}\right)^2\lambda \tag{2.7}$$

When trying to record such small objects like tracer particles, the CCDs or CMOS pixel size needs to be accounted. If the pixel size exceeds the size of the individual particle image, then the entire position of the particle, including its light intensity distribution, will be lost as depicted in Figure 2.5 (left). Particle images sizes should be between 2–4 pixels to keep the effects of peak-locking error to a minimal level [9] [10].

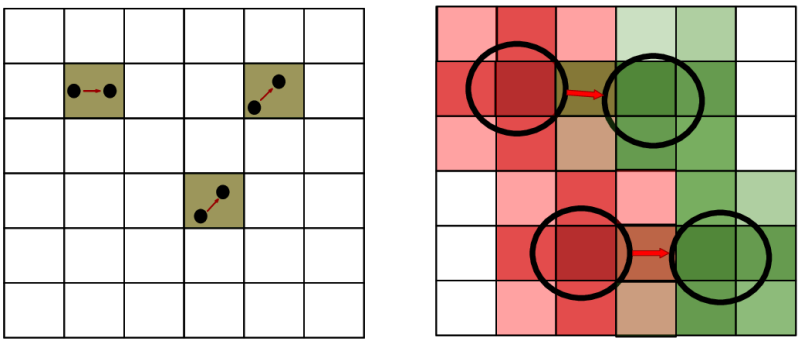

(a) under-resolved particle (b) well-resolved particle

Figure 2.5: Digital imaging of small particles. Source: [11].

When two consecutive recordings have been acquired, the images are divided into small sections, called *interrogation windows*, that contain a statistically significant number of tracer particles. The interrogation windows from each image frame are cross-correlated with each other, pixel by pixel. This process returns the discrete cross-correlation map, whose highest signal peak

position relative to the origin indicates the average particle displacement. An image's smaller scale is a pixel, which means that the resulting particle displacement will be given as an integer number of pixels that the particle has moved in time. Therefore, in order to obtain the particle image displacement with sub-pixel precision, the correlation peak needs to be interpolated around its maximum. This is commonly done by approximating the cross-correlation curve to a Gaussian curve. By repeating the process to the whole image area, the velocity field over the recorded flow domain is obtained.

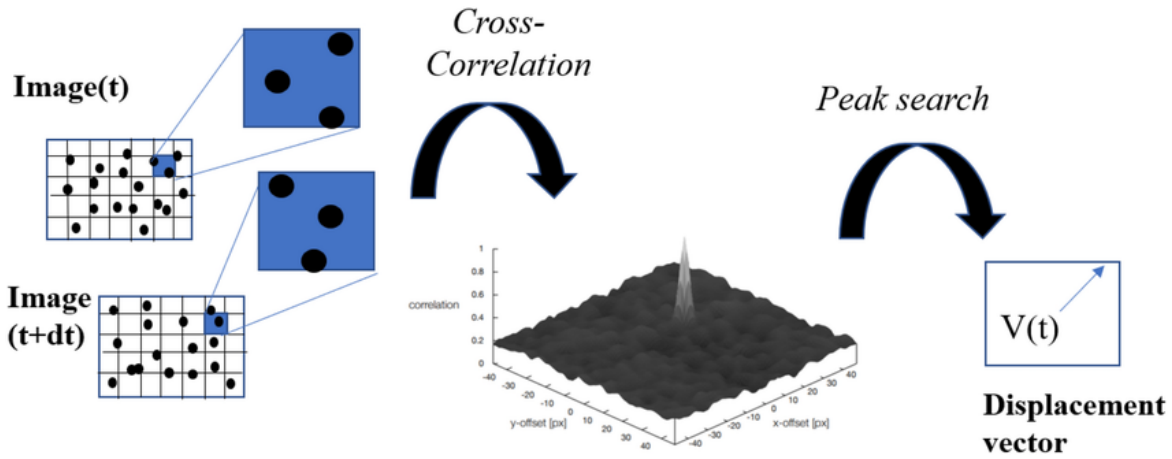

Figure 2.6: Sketch of cross-correlation and peak search for velocity vector determination. Source: [12].

PIV presents several advantages that makes the technique very appealing. It is a non-intrusive velocity measurement, which means that no objects that could disturb the mechanical properties of the flow domain are introduced into the fluid. This enables a more accurate measurement, compared to other techniques like hot-wire velocimetry (HWA) or pressure probes. PIV is a whole field technique, therefore it does not provide point information within the flow domain, but allows capturing instantaneous and time-resolved information on the whole flow field domain.

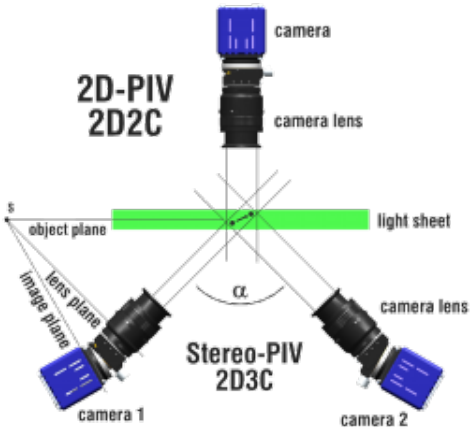

Figure 2.7: Planar and Stereo-PIV set-ups. Source: [13].

Planar PIV (2D-2C) was introduced in the decade of the 80's as an innovative and cutting-edge technology to measure instantaneous two velocity components (u and v) within the plane domain (x, y). Since then it has been used in a wide variety of areas. Guida et al. investigated the effects of the azimuthal position of the measurement plane in a fully baffled vessel agitated by a pitched

blade turbine [14]. Terra et al. performed 2D PIV experiments on the flow in the wake of a cylinder in cross flow [15]. Brito et al. used the technique to get the flow field on the middle section of a coiled copper pipe [16]. An analysis of multiple air jets impinging on a moving flat plate was carried out by Barbosa et al. [17]. In the field of sports aerodynamics, Jux et al. [18] analysed three different bike wheel tires in cross wind conditions with planar PIV measurements in the TU Delft Open Jet Facility.

The measurement of the instantaneous velocity vector fields with the three velocity components (u, v and w) within a two-dimensional plane domain (x, y) is possible thanks to planar stereoscopic PIV (2D-3C). By introducing a second camera in an angle (see setup in Figure 2.7) and considering that the laser sheet has a certain thickness, the third velocity component can be retrieved as the out-of-plane velocity component [19]. Spoelstra et al. used large-scale stereoscopic PIV measurements on a plane crossed by a cyclist to analyse the flow in action (Ring of Fire technique) [20]. This technique has also been implemented on the health field, for instance on the analysis of the rupture risk of intracranial aneurysms [21]. [22]. The wind energy field also makes use of this technique; some studies involve the analysis of the wake deflection of a vertical axis wind turbine (VAWT) [23] or rotor–wake and wake–wake interaction of VAWTs [24].

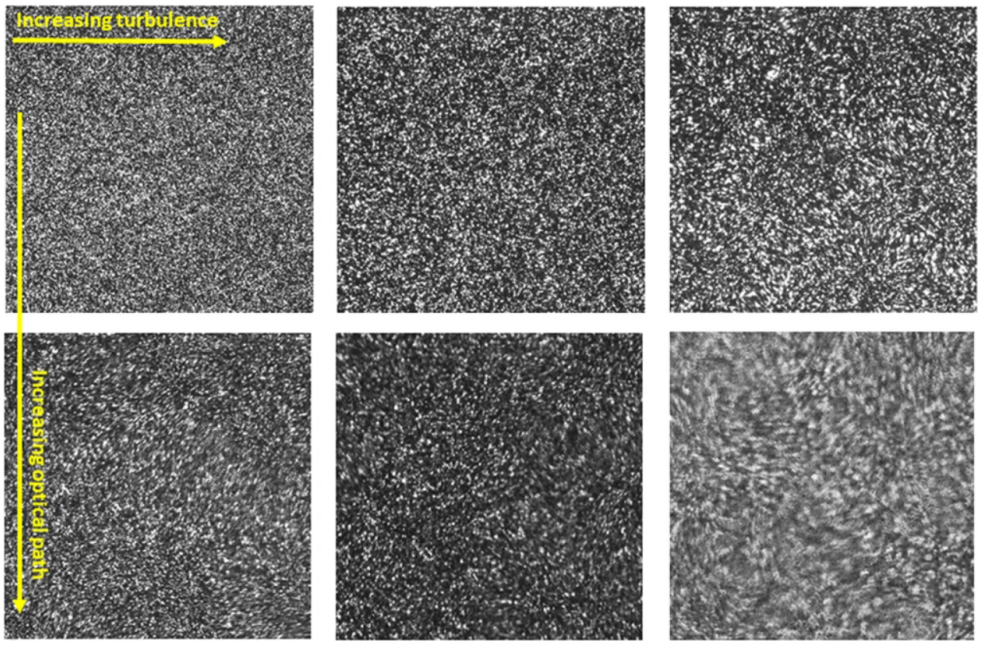

Figure 2.8: Planar PIV raw images. Source: [25].

2.1.1. 3D Volumetric PIV

Planar and stereoscopic PIV provide information about the velocity field in a 2D space. However, turbulent flows are naturally 3D. Therefore, in order to be able to better understand them, these should be resolved in all three dimensions, which is not possible with either planar or stereo PIV. An approach that can be followed to retrieve information about a volume domain is multiplane Stereo-PIV, which consists on performing multiple 2D planes measurements and, by comparing them together, interpret the velocity field in a 3D space. Kahler invesitaged fully developed turbulent boundary layer flow along a flat plate [26] and turbulent mixing in wall bounded flows [27] by using multiplane Stereo-PIV. On the medical side, the flow within a transparent model of a human lung is also studied by Schröder et al. [28] by means of this technique. However, this process has limitations with its application. On one side, the 3D velocity field extracted from

the different 2D planes is not instantaneous since they are not acquired at the same time. In consequence, this technique would only suitable if the acquisition time between planes is smaller than the minimum time scale of the flow. Also, high-speed cameras would be needed, which increases the cost of experiment.

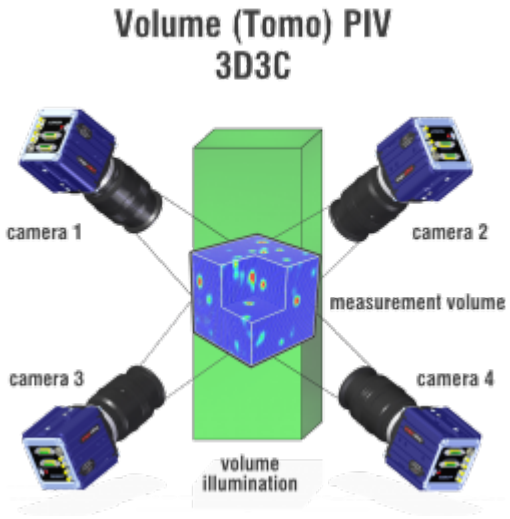

Figure 2.9: Tomographic-PIV set-up. Source: [13].

In 2006, Elsinga et al. [29] introduced Tomographic Particle Image Velocimetry (Tomographic-PIV or Tomo-PIV) which, thanks to tomographic reconstruction techniques using multiple views, allowing to capture instantaneous three-dimensional velocity vector field volume (3D-3C). The working principle is similar to planar PIV: the velocity field is obtained by the displacement of the particles from two subsequent images. However, there are some notable differences between both approaches. Figure 2.9 depicts a simplified set-up for Tomo-PIV.

The imaging of the tracer particles in a Tomographic-PIV experiment is conducted at illumination intensity typically an order of magnitude smaller than that of planar PIV due to the expansion of the laser beam over a large cross section. The problem is further exacerbated by the small optical aperture (high f-number) of the imaging system, needed to ensure focused particles across the whole measurement depth. As a result, the peak intensity of particle images decreases by almost an order of magnitude when the volume depth is doubled. As a consequence, the intensity counts of particle images rarely exceed a few hundreds.

The same tracer particles described in Section 2.1 can also be used for volumetric PIV. However, in 2015, neutrally buoyant Helium-filled Soap Bubbless (HFSB) were introduced as an alternative tracer particles to overcome the limitations of light scattering in larger-scale cases [30]. These consist of soap bubbles filled with a gas that has a smaller density than air (typically helium). By compensating the weight of the soap with the volume of the helium inside, the neutral buoyancy condition can be obtained [31]. Thanks to their diameter of 0.3 mm and a response time of less than 15 μs scatter 10000x more light than μm-particles, it makes them suitable for large scale PIV/PTV experiments in the lower subsonic regime.

The measurement flow domain is not a plane, but a volume. Therefore, to illuminate this volume, the thickness of the laser light sheet is expanded thanks to a beam expander. The intensity of the light is inversely proportional to the thickness of the measurement volume since, as the laser expands it is less concentrated, making its intensity weaker. This is a limiting factor on the range of measurement volume size able to illuminate and the light intensity scattered by the tracer particles. The latter limitation is also related to the recording device characteristics. In order

to track the particles, all of them need to be in focus across the whole depth of the measurement domain. To achieve this, the focal depth of all the cameras should be significantly long, which normally is possible by increasing the numerical aperture of the objectives until the focal depth matches the thickness of the volume illuminated. This can also be achieved by using lens-tilt adapters that align the focal plane with the mid plane of the measurement volume (Scheimpflug condition [29] [32]).

2.1.1.1. Tomographic reconstruction

To retrieve the instantaneous measurement of all three velocity components in a complete 3D measurement volume (3D-3C), at least three cameras pointing with different angles towards the measurement volume are required. The more cameras are used, the more additional information on the particles from different angles, which increases reconstruction accuracy. For a 4-camera set-up, as shown in Figure 2.10, four pairs of images with different views are obtained. After the image recordings are acquired, the tomographic reconstruction step for both time steps takes place, which is basically a mapping function between image and object spaces.

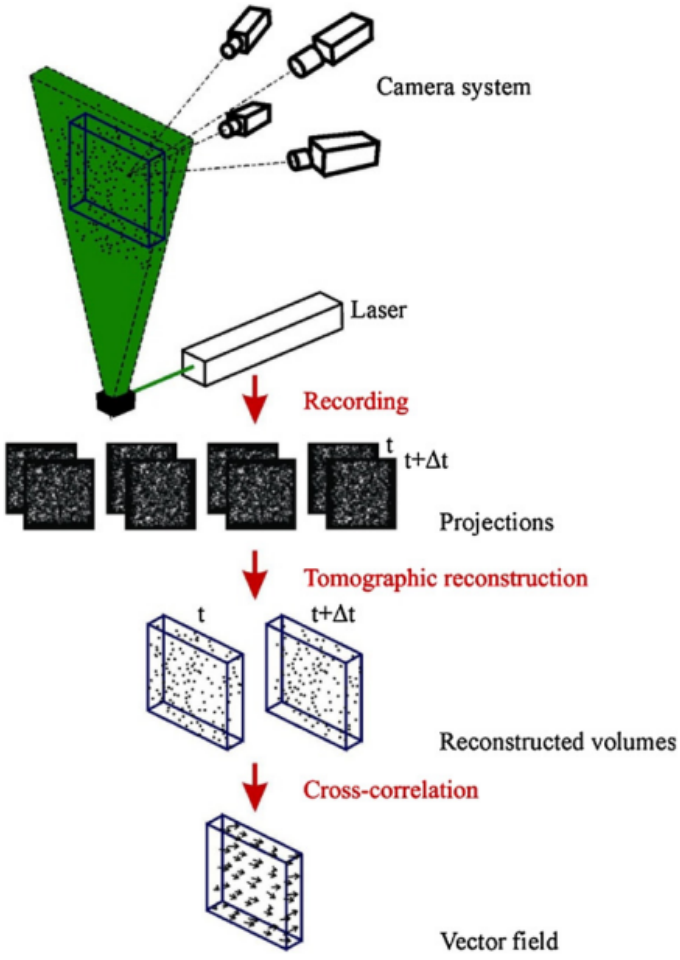

Figure 2.10: Tomographic-PIV flow-chart. Source: [29].

The measurement volume containing the particle distribution (the object) is discretized as a 3D array of cubic voxel elements with their intensity value. Voxel is short for *volume pixel*, a 3D equivalent of a 2D pixel. The gray values of each pixel from every camera are projected into the voxel space by a weighting matrix where the different values are multiplied with one another. The

2D images are reconstructed into the 3D particle distribution volume by means of a tomographic reconstruction algorithm (Multiplicative Algebraic Reconstruction Technique, MART) as a 3D light intensity distribution for each voxel [33]. Then, the three-dimensional cross-correlation is performed, allowing to retrieve the 3D velocity field within the measurement volume.

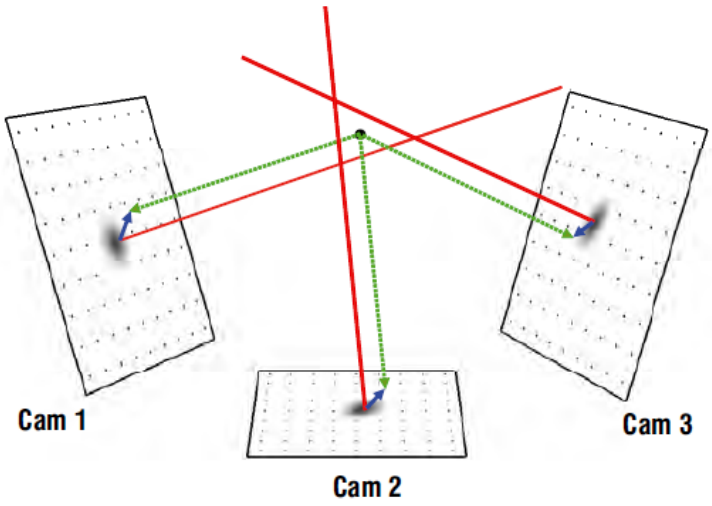

Figure 2.11: Volume self-calibration. Red line: original projection; black dot: assumed correct particle position; dotted green line: corrected back-projection; solid blue arrow: disparity vector. Source: [34].

In order to mitigate reconstruction inaccuracies, a 3D self-calibration [35] of all the cameras must be done. Each camera records images of a reference calibration plate from different views. Then, the particles or calibration points are detected and their triangulation in 3D space is performed. Due to inaccuracies in the calibration function, particles in the measurement volume are imaged at slightly shifted positions in the camera images - see Figure 2.11. By averaging these variations across the set of particles within a local sub-volume, 3D disparity maps are constructed, leading to corrections in the calibration function. This calibration procedure establishes the relation between image coordinates and the object or physical space.

Figure 2.12 shows the resulting disparity maps of a calibration procedure for a four-camera tomographic PIV setup, where the top-left shows the reference calibration pattern. Top-right and bottom-left show acceptable calibration results with the peaks of the ”particles” in the center of each quadrant, whereas the bottom-right presents displaced peaks, probably due to camera displacement after calibration. The quality of this procedure will affect the accuracy in the reconstruction step.

One of the main issues of this technique is the appearance of *ghost particles*. These are a number of particles that appear with respect to those actually present in the illuminated measurement volume. This phenomenon increases with higher particle density, the particle diameter and the length of the line of sight in the volume. An increased particle density produces a larger percentage of ghost particles, thus decreasing the reconstruction quality. However, at the same time, a large number of particles allows a higher spatial sampling rate of the flow, returning a potentially higher spatial resolution. Hence, a high particle density is preferred but always considering the effect of ghost particles. A way to minimize this effect is to increase the number of cameras; this allows capturing the measurement domain from different views and therefore, improving the triangulation of real particle images in the reconstruction step. Another limitation is the high processing time required for the 3D reconstruction and the velocity field computation, which supposes a high computational cost.

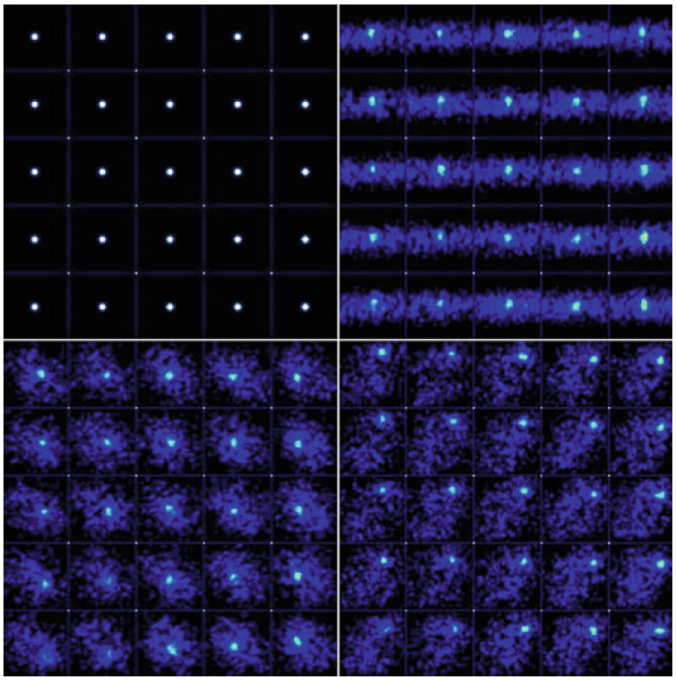

Figure 2.12: Disparity maps for a four-camera tomographic system obtained with the 3D self-calibration technique. Source: [3].

2.1.1.2. Shake-the-Box

Shake-the-Box (STB) is a 3D Lagrangian Particle Tracking Velocimetry (PTV) algorithm [36]. It combines the calibration methods (volume self-calibration and OTF, Optical Transfer Function, calibration) of Tomographic-PIV and the iterative triangulation and the image matching (*shaking*) by Iterative Particle Reconstruction technique (IPR). The process can be divided into the following three stages:

- *Initialization*. First, an initial prediction of the particles locations over a few recordings (typically the first four time-steps) is done by means of the particle-based *Iterative Particle Reconstruction technique* (IPR) [37], whose objective is to pair particles between frames. The identified particles are considered *particle candidates* if such particles for which a track can be identified are considered as true, reconstructed particles. The rest of the particles (untracked) are potential ghost particles, particles that do not exist but can appear due to reconstruction. After identifying the particles, the corresponding trajectories are extracted from the distributions of particle candidates. Consecutive frames are checked and matches are to be found by applying a search radius around either the particle position or a predictor location.
- *Convergence phase*. A prediction for the particles position at the next time-step is done from the initial particles positions obtained in the previous initialization stage. The particle tracks predicted position is slightly modified in a process referred to as *shaking* with the aim to correct for possible deviations in particle positions between the predicted projected image and the original image. By doing the residual of the predicted and original images, new particle positions appear and these are stored as candidate particles. This process if repeated for every time-step. In the case that these candidate particles appear in four consecutive time-steps, they are considered as *real particles* and are kept in the residual, otherwise they are removed. This process is represented in Figure 2.13.
- *Converged phase*. The algorithm needs some time-steps where the number of particles

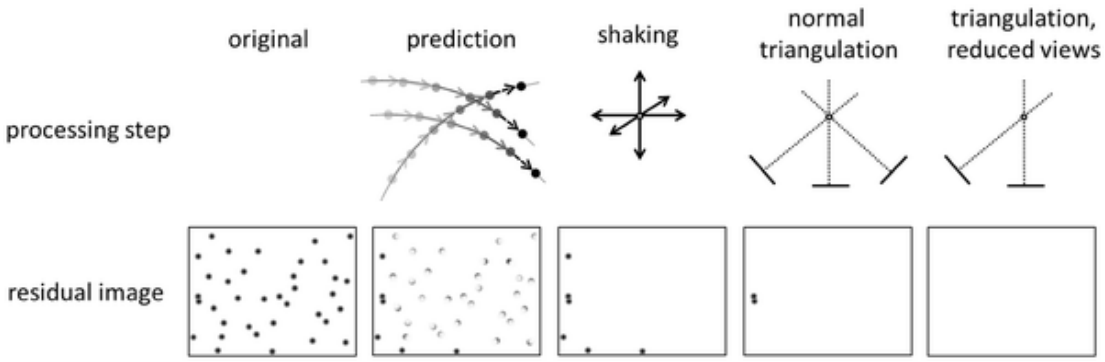

Figure 2.13: Schematic representation of the Shake-The-Box (STB) for a single time-step. Source: [36].

between time-steps does not change significantly, therefore the process reaches its end, its *converged* state. In this stage, these are mostly new particles that have entered the measurement domain within the last four time-steps.

By tracking the position of each particle tracer at a specific time step Δt, STB allows evaluating the displacement of each of these particles captured within the measurement volume. This approach allows fast processing of three-dimensional data with high particle concentrations, while capturing the vast majority of true particles and minimizing ghost particles.

The Shake-the-Box algorithm allows its implementation in different modes:

- **Time-Resolved Shake-the-Box (TR-STB).** This is the STB algorithm described above.
- **Double-Frame Shake-the-Box (DF-STB).** In some cases, time-resolved measurements cannot be performed due to hardware limitations, especially when the flow of interest is too fast for the available cameras or illumination characteristics. In this situation, an alternative approach is Double-Frame Shake-the-Box (DF-STB). In this mode, two sets of double-frame images are recorded at a pulse separation of different duration (one shorter and one longer). Saredi et al. [38] proposed a method in which the set of images measured with the shorter time separation is employed to generate a robust displacement predictor that is used on the longer time separation to extend and complement the measurement.

2.1.1.2.1 Binning

The data obtained from Shake-the-Box shows the velocity of particles tracked over time as scattered data over the measurement domain, commonly called unstructured data. In order to simplify the results and facilitate their interpretation, a *binning* step can performed to transform the unstructured velocity information to a structured one.

Binning is a data reduction technique used with the aim to reduce the size and complexity of 3D volumetric flow measurements datasets. It is based on grouping neighboring velocity vectors within specific previously defined volumes or cells called *bins*. After this process, a single velocity value is assigned to each of these bins, resulting in a structured velocity field. This is particularly advantageous in the case of high-resolution images that contain a big amount of individual particle track data or when there is a lot of noise or outliers in the data. With the binning, an averaging of the flow field is being done, which helps reducing the noise in the data. However, this data reduction approach is case-dependent, so its processing parameters need to be carefully selected to obtain the desired data with minimally losing the most important features of the flow.

2.1.1.3. Robotic Volumetric PIV

Robotic PIV consists of a robotic arm fitted with the cameras and the laser light source, all aligned in a compact housing. Figure 2.14 shows an example of how this technique is used to measure the flow around a cyclist [39]. The robot arms allows a more flexible data acquisition since if a new volume wants to be measured, just the robot sequence needs to be changed. This reduces the complexity of the experimental setup in the test section that can suppose Tomographic PIV (as seen in Figure 2.9).

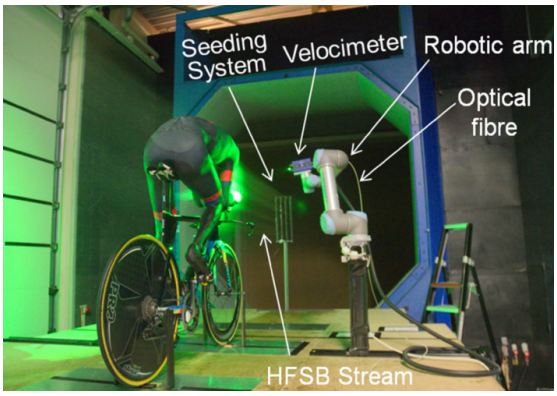

Figure 2.14: Robotic PIV set-up in a wind tunnel. Source: [39].

The Robotic PIV configuration is similar to Tomographic PIV (recall Figure 2.9), but as previously discussed, instead of having a fixed recording and illumination systems, these are located within a robot arm that facilitates the recording of measurement volumes from different angles. Multiple 3D velocity measurements are performed with Coaxial Volumetric Velocimetry (CVV) setup [40] as shown in Figure 2.15. Since the cameras and the laser are aligned, the proper volume measurement acquisition is guaranteed without the need to manually align them, which could lead to human errors and a compromise to the measurement accuracy.

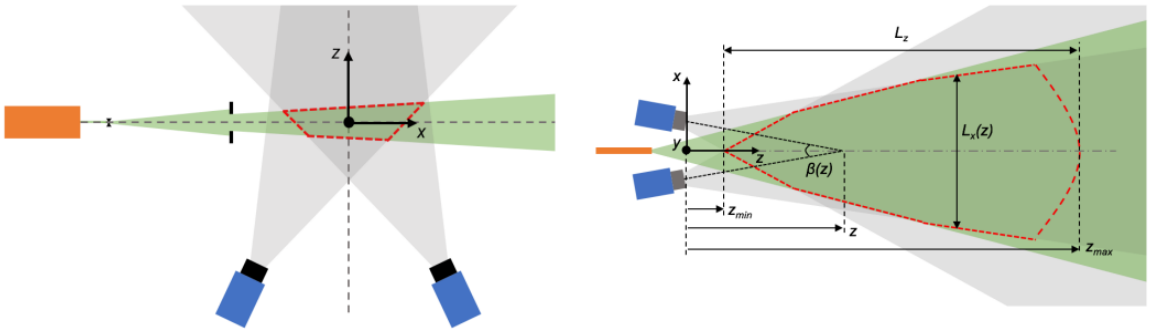

Figure 2.15: (Left) Tomographic PIV setup and (Right) Coaxial velocimeter (CVV) setup. Where cameras (blue), field of view (grey), laser illumination (green) and optical fiber (orange). Source: [40].

Furthermore, knowing the position of the robot base with respect to the domain reference frame and the robot head CVV position with respect to the robot base, it is possible to transform the particles reconstructed (e.g. by STB) into the main reference frame and, hence obtain the 3D velocity field in the physical space.

2.2. Reflection Treatment Techniques

As exposed in Section 2.1, the basic fundamental requirement of PIV is to be able to track particles. When these can be properly tracked, their displacement can be determined, and thus velocity field can be retrieved. At the moment the particles can no longer be traced or are not visible in the images, the accuracy in the velocity field calculation is compromised.

Some of the best practices to record PIV images for a proper particle-tracking would be: tracer particles big enough to be visible, camera characteristics such that all particles are in focus, optimum seeding density (sufficient particles in the field of view but avoid tracers to overlap) and high contrast between background and particles.

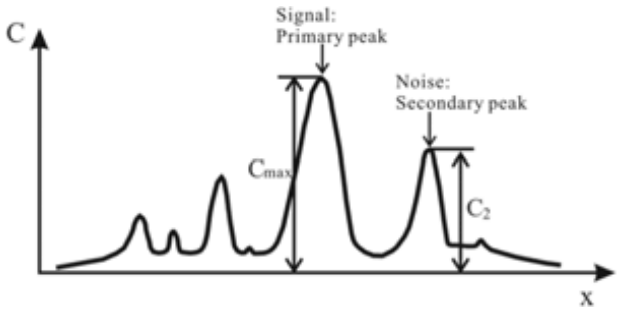

Figure 2.16: Signal-to-Noise ratio 1D representation, where *primary peak* is *peak 1* and *secondary peak* is *peak 2*. Source: [41].

The latter plays a significantly important role to accurately detect and track particles and can be quantified by the Signal-to-Noise ratio (SNR), which is defined as:

$$SNR = \frac{peak1}{peak2} \tag{2.8}$$

peak 1 is the highest peak of the cross-correlation, representing where the particles coincide from one frame to the other. The cross-correlation also shows additional peaks due to noise and the correlation of non-paired particles. Hence, *peak 2* is the second highest peak and can be interpreted as the background noise of the image. A high SNR yields a situation with dark particles on a brighter background, which allows distinguishing more easily particles from the background. Otherwise, in the case of low values of SNR, the particle can no longer be distinguished from the background, thus their tracking would not be possible.

An important phenomenon that can lead to low SNR values (and even loss of particles) is light reflections. These are high intensity (typically one order of magnitude higher than particle intensity) areas that appear due to the laser light impinging on a surface; when the angle of incidence is close to 90 degrees, this light reflection becomes more intense. Figures 2.17 and 2.18 serve as examples of the effect of these reflections on PIV images and the capability to detect and track particles. Left images clearly show that the particles in the surroundings of the surface will almost be untraceable, whereas right images mitigate these reflections but still keep the presence of the tracers.

These light reflections imply a big detriment when retrieving the velocity field of the measurement domain. When the cross-correlation is performed, reflections can appear as a peak that would affect the detection of the displacement peak. By human eye, it is clear that a reflection is a reflection when it appears in a PIV raw image. However, removing them manually would be a tedious job and it is not feasible with the high amount of data acquired in a PIV experiment (specially with Volumetric PIV). For this reason, it is particularly important to have an efficient

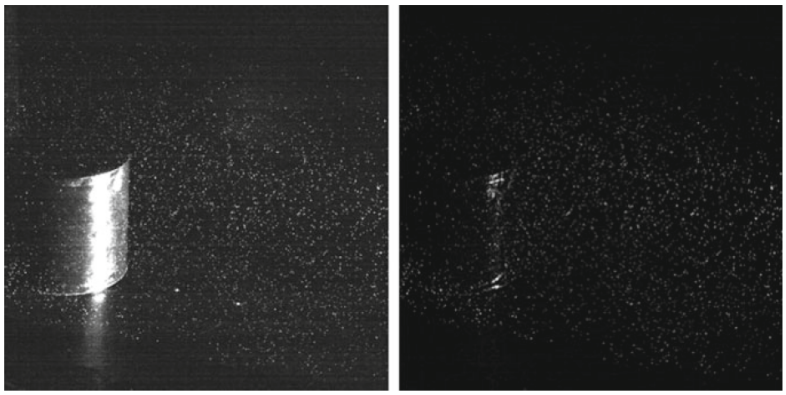

Figure 2.17: Flow in the wake of a cylinder obstacle raw image taken by tomographic PIV (left) and pre-processed image (right). Source: [3].

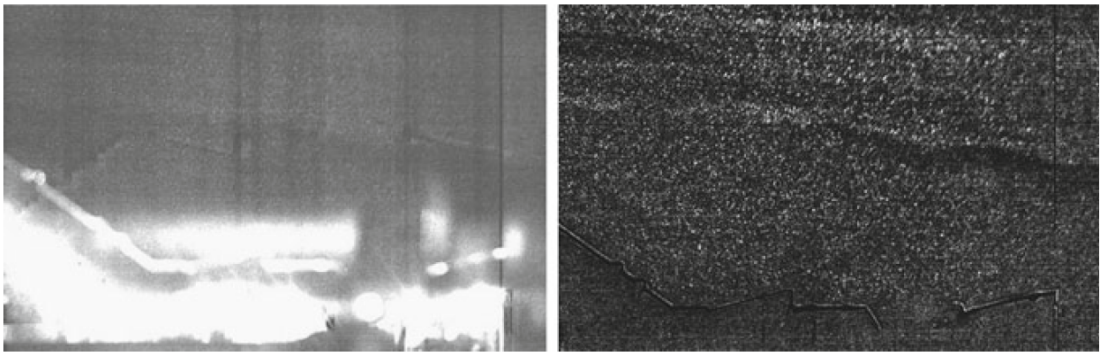

Figure 2.18: (Left) Raw image of Ariane 5 launcher and (right) pre-processed image with temporal high-pass filter. Source: [3] [42].

procedure to completely remove them or at least minimize their intensity with the aim to have a high SNR for particles to be visible and detectable. In the last few years, there have been a few studies to best define a procedure to mitigate this phenomenon.

The following sections aim to serve as a review of the current methodologies for reflections detection and their mitigation in different stages of the PIV pipeline.

2.2.1. Reflection treatment methodologies in literature

As described in the previous sections, the PIV pipeline is composed of different steps. There have been a few studies on reflection detection and mitigation in the different stages of PIV pipeline.

The first simplest option that one could think of to avoid light reflections on the surface of the model would be to modify the surfaces where the cameras are pointing. One way to do it is to apply a treatment on the surface of the model. A common strategy is to use matt black paint to cover the model which makes most of the light to be absorbed instead of being reflected. Another common approach is to either paint the model with a fluorescent coating (e.g. Rhodamine) ([43], [44]) or add a fluorescent dye to the particles ([45], [46], [47], [48]). The light reflected by either the surface or the particles will shift with a different wavelength than the laser light and, by adding a band-pass optical filter on the camera, the shifted wavelength will be filtered and only the laser light will be picked up. Another option would be to change the camera view angles with the aim to avoid light reflections impinging perpendicular to the camera; get a camera view that captures the surface minimizing the associated reflections ([49], [50]). Nevertheless, in some cases, applying

the latter techniques is not trivial as complex geometries can suppose a barrier to these set-up changes for reflection mitigation. For this reason, alternative approaches need to be taken on the PIV stages concerning data processing.

A statistical model by means of a median detector was introduced by Westerweel [51] to find the occurrence of spurious vectors in PIV data. Westerweel and Scarano [52] did a follow-up on the latter approach by establishing a normalized median with respect to a robust estimate of the local variation of the velocity, which is commonly called *universal outlier detection*. Wang et al [53] applied proper orthogonal decomposition (POD) to detect and replace outliers by dynamically approximating the original pure velocity field. [54]. A way to classify the flow field can be classified depending on the type of outliers was set up by Tang et al [55] by using the penalized least-squares (PLS): PIV vector fields containing scattered outliers are detected and corrected using higher-order differentials, while lower-order differentials are used for the flows with clustered outliers. Recently, Saredi et al [56] proposed velocity field outlier detection based on the turbulent kinetic energy (TKE) transport equation; the ratio between local advection and production terms of the TKE on a streamline determines whether the data is admissible or not.

Despite the previous reflection detection approaches, the most commonly used methodologies involve pre-processing of raw images. This allows the removal of the pixel intensity corresponding to light reflections in a stage previous to the velocity field computation. A widespread approach is background removal of PIV images, which consists of generating a background (or reference) image and then subtracting it from the raw image. This reference image can be obtained by means of several ways: recording an image without particles, obtaining the average or local minimum (Minimum Subtraction) of a set of recordings ([57], [58]), for instance. Willert [8] proposed using a high-pass filter by subtracting a low-pass-filtered image from the original (typically a 7x7 pixel smoothing kernel) and then the resulting image is a smoothed with a 2 × 2 pixel kernel filter. This showed a good performance in bringing most particle images to the same intensity level. LaVision also introduced an option that subtracts the sliding temporal minimum or sliding average intensity of each pixel over multiple time steps [59]. Honkanen and Nobach [60] proposed a double-frame image pre-processing based on subtracting the second frame from the first frame. Their main assumption is that what is kept still is assumed to be a source of bias and is removed; in this case, light reflections would be included in this source of bias as they generally stay still for more than one frame. However, this procedure is limited by the fact that the particle displacements should be more than one particle image width, not to consider also the particle images as a source of bias and subtract them. As a continuation of the latter, Mejia-Alvarez and Christensen [61] introduced a modification by computing a local-median normalization of the intensity with respect to the difference between sliding median and minimum intensities. [62]. Wang et al [63] proposed a cut-off filter based on the ratio between the mean gray-scale intensity map and the original image (called *ratio cut method*). This approach is shown to be unsuitable for time-varying reflections with strong intensity since their variance affects the average intensity map. Also, a non-dimensional threshold has to be manually set, which does not have a universal value yet, hence a study on which value to set should be done every time this approach is applied.

There are three image pre-processing approaches to highlight that show good performance on minimizing light reflections on PIV raw images: (1) temporal high-pass (Butterworth) filtering, (2) background image via Proper Orthogonal Decomposition (POD) and (3) background image generation via anisotropic diffusion. These are presented more in depth in the following subsections.

2.2.1.1. Temporal high pass filter

The use of a temporal high-pass filter for the elimination of both steady and unsteady reflections was proposed by Sciacchitano and Scarano [64]. The intensity signal in time for a specific pixel is shown in Figure 2.19. Here, reflections and particles can be clearly distinguished from each other: in the case of a reflection, its intensity appears as a much higher peak and longer duration of high intensity, which represents low-frequency content. On the other side, particles show as short pulses of a much lower intensity, hence would be contained as high frequencies in the frequency domain.

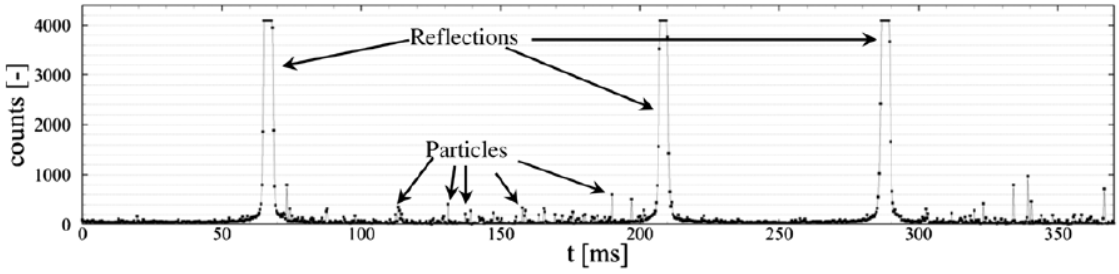

Figure 2.19: Intensity signal at a pixel location in time. Source: [64].

The proposed method consists of the decomposition of the pixel intensity signal in the frequency domain at individual pixel locations by means of Fourier transform. Then, assuming that high frequencies represent the transit of seeding particles, while low frequency is unwanted reflections. In order to filter the low frequencies associated with light reflection, a high-pass filter (HPF) is applied on the light intensity signal. Specifically, a Butterworth HPF is used for this purpose.

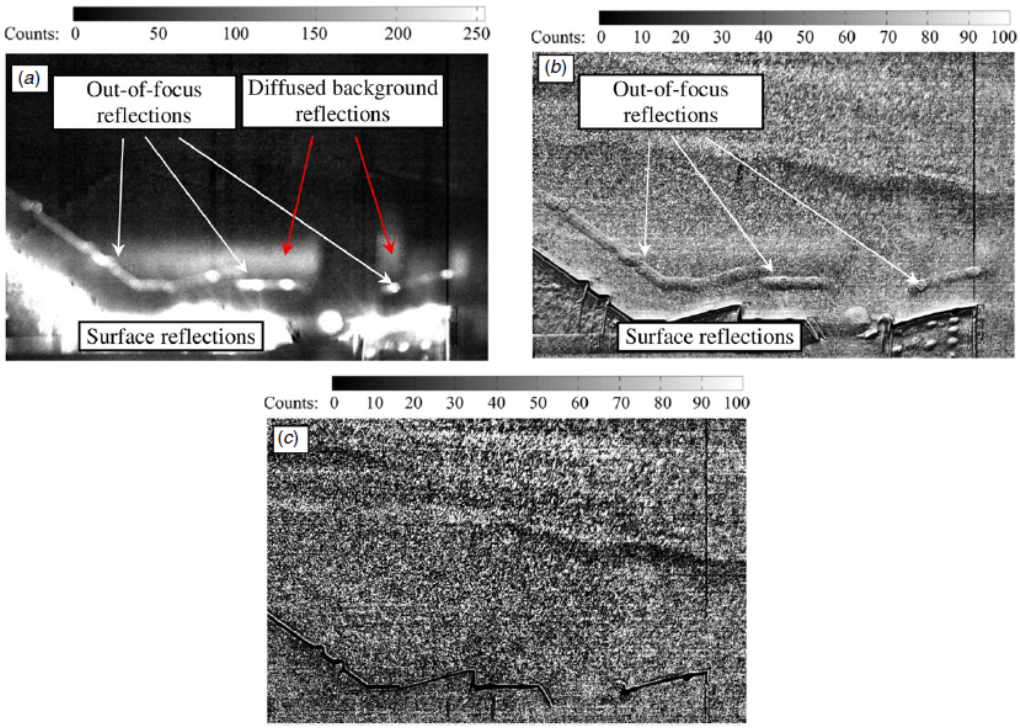

Figure 2.20: ARIANE V after-body (a) PIV raw image, (b) minimum intensity subtraction (c) Butterworth HPF with cut-off frequency 30% of Nyquist frequency. Source: [64].

The main advantage of this approach is that it can be applied on both steady and unsteady reflections. However, there are a few limitations to bear in mind. In the case of the pixel intensity increase rapidly from 0 to maximum intensity value corresponding to a reflection, the use of the temporal high pass filter would strongly attenuate the reflection, but will not completely remove it. It also relies on the assumption that the separation of timescales between particle and object movements should be long enough. Therefore, when objects move at similar speed as surrounding particles, this approach would not be suitable.

2.2.1.2. POD-based background removal

Mendez et al. [65] proposed using Proper Orthogonal Decomposition (POD) of a sequence of PIV images with the aim of generating a Reduced Order Model (ROM) that only maintains the PIV particles. The idea would be to apply a POD filter that is able to automatically detect and remove the minimal number of modes that represent the background noise.

Figure 2.21 presents a PIV raw image and a comparison of different methodologies to remove reflections. These results show that POD performs quite well compared to the other techniques presented. This approach is independent of temporal resolution of recording sequence, and the sharpness or intensity of the background noise.

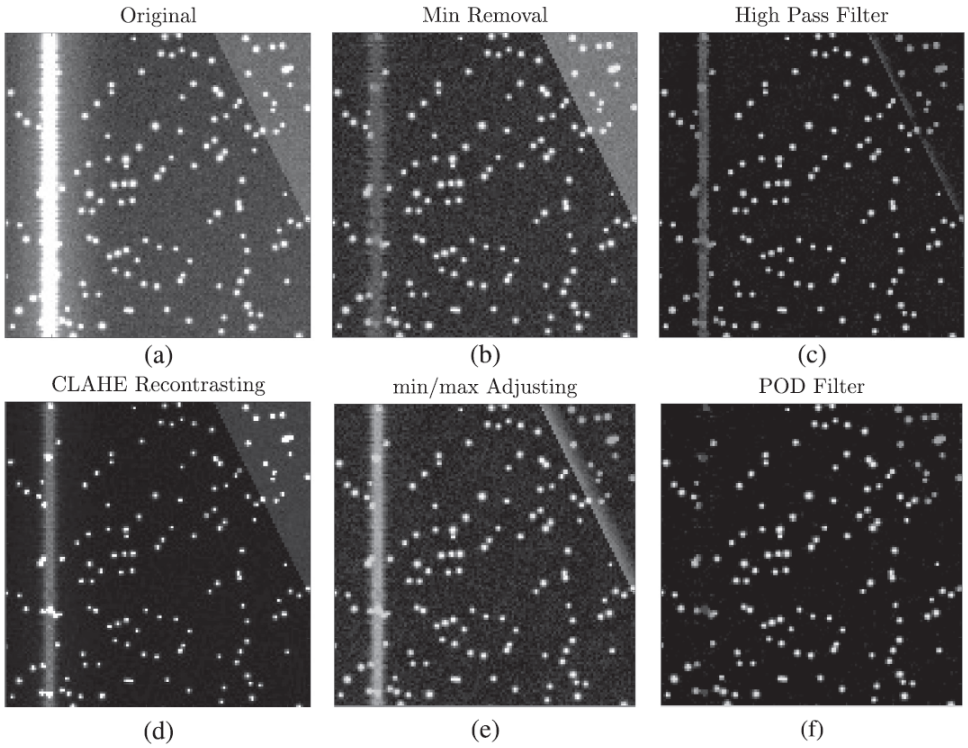

Figure 2.21: (a) Raw image, (b) minimum background substraction, (c) Butterworth high pass filter, (d) CLAHE recontrasting, (e) mininum/maximum adjusting and (f) POD filter approach. Source: [65].

This method requires a large dataset of images for the POD to converge. It also need that the reflection stays in the same or approximate pixel locations for several number of images, otherwise the reflection is at risk of being confused with particles that come and go and will not be removed. Its dependence to threshold selection for ROM reconstruction and its high computational cost also constitute as additional limitations.

2.2.1.3. Anisotropic diffusion-based background removal

Adatrao and Sciacchitano [66] proposed the generation of a background image by anisotropic diffusion of the intensity distribution of the raw image. The main idea is to consider that diffusion occurs along the edges and not across them, leaving large areas of high intensity, which represent reflections, in the background image. Then, this background image is used to be subtracted from the raw image, resulting in a pre-processed image with no reflections but with the presence of particles.

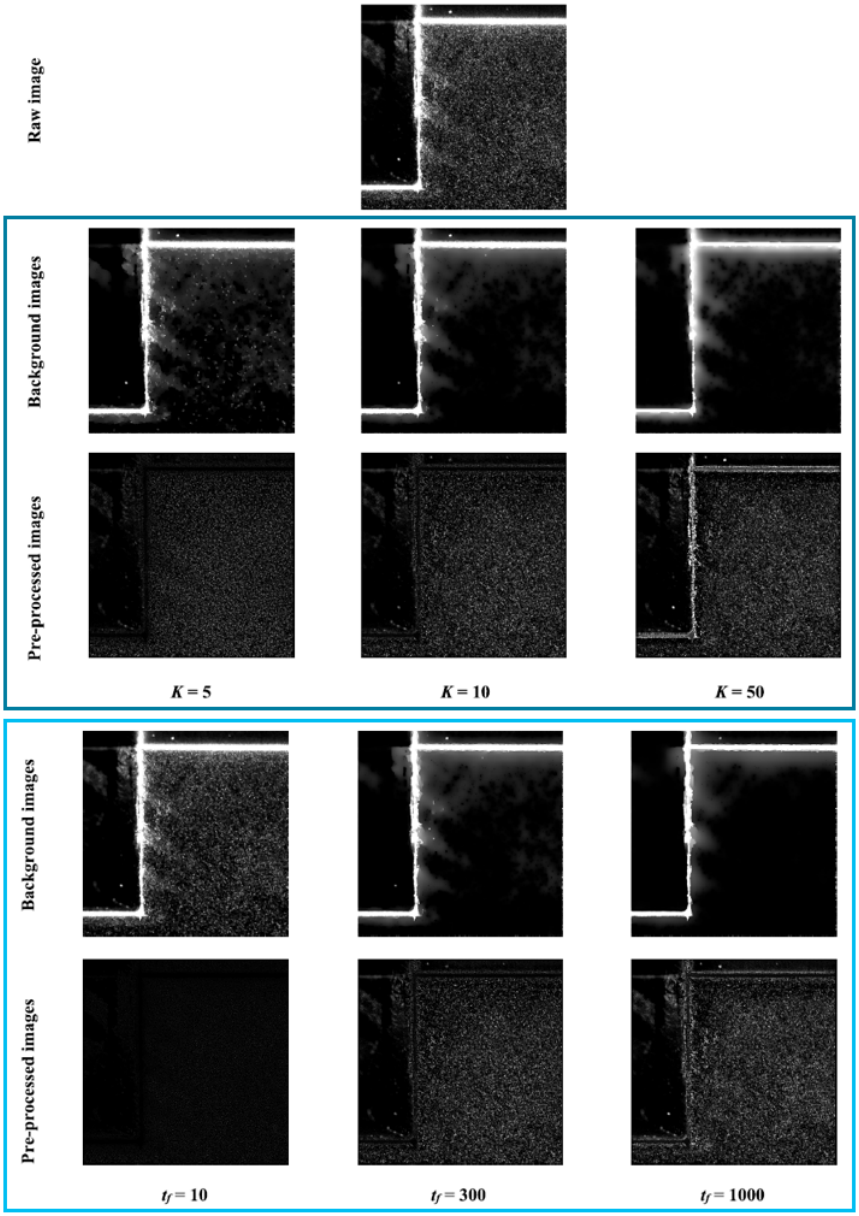

Figure 2.22: (Top) Raw image with comparison of background and pre-processed images obtained by anisotropic diffusion method with (middle two rows) t_f = 300 and different threshold numbers (K = 5, 10, 50) and (bottom two rows) K = 10 after different numbers of iterations (t_f = 10, 300, 1000). Source: [66].

The background image intensity can be obtained as the solution $I(x, y, t)$ of the following diffusion equation:

$$\frac{\partial I}{\partial t} = \nabla \cdot [c(x, y, t)\nabla I] = c(x, y, t)\Delta I + \nabla c \cdot \nabla I \tag{2.9}$$

where c is the diffusion coefficient as a function of the magnitude of the intensity gradient $|\nabla I|$ and g a monotonic function:

$$c(x,y,t) = g\left[\nabla I(x,y,t), I_n\right] \qquad g(x,y,t) = \frac{1}{1+\left(\frac{|\nabla I|}{K \cdot I_n}\right)^2} \tag{2.10}$$

The diffusion coefficient is calculated considering both the magnitude of the intensity gradient $|\nabla I|$ and the local normalized intensity I_n. This allows distinguishing between reflections (bigger areas of bright pixels) which are considered to have small values of local normalized intensity I_n compared to $|\nabla I|$ and small bright spots (i.e. particle images) with large values of I_n compared to $|\nabla I|$. To be able to solve Equation (2.9), the threshold parameter K and the number of iterations t_f must be determined. It is found that higher values of K lead to isotropic diffusion, hence smoothing the reflections and t_f defines the number of pixels that will be taken into account in the diffusion process. The authors advise to perform a study on both these parameters before applying this approach.

This methodology is suitable for removing reflections single-frame PIV images since it only takes into account the intensity distribution of the recording. Therefore, it can be applied to either steady and unsteady reflections (e.g. propellers, pitching airfoil). Conversely, an important limitation of this procedure is its dependence on good selection of the threshold parameter K and number of iterations t_f, as shown in Figure 2.22. A wrong calculation of these could lead to an unsuccessful background removal.

2.3. Fourier Analysis in image processing

2.3.1. 1D Fourier transform

Fourier analysis is named after Jean Baptiste Joseph Fourier (1768-1830), a French mathematician and physicist. He made the following claim: *any continuous periodic signal can be represented as the sum of sines and/or cosines of different frequencies, each multiplied by different coefficients* [67]. This sum is what is commonly known as *Fourier series*. Non-periodic functions can also be expressed as the integral of sines and/or cosines multiplied by weighting function, process that is called *Fourier transform*. Both the Fourier series and transform have a common key characteristic: they can be reconstructed completely by means of an inverse process without any loss of information. This allows the user to work in the Fourier or frequency domain and then come back to the original domain without losing any information about the function.

The *Fourier transform* decomposes any function in the temporal or spatial domain into a sum of sinusoidal basis functions in the frequency domain. Each of these basis functions is a complex exponential of a different frequency in Hertz (Hz) or the number of cycles per second.

- In the time domain t:

$$G(f) = \int_{-\infty}^{\infty} g(t) e^{-j2\pi ft} dt \tag{2.11}$$

- In the spatial domain x:

$$G(k_x) = \int_{-\infty}^{\infty} g(x) e^{-j2\pi k_x x} dx \tag{2.12}$$

where $g(t)$ and $g(x)$ are continuous temporal and spatial signals, respectively with f being the frequency in the time domain (Hz or cycles per second) and k_x the frequency in space domain (cycles per unit of space); the two latter are continuous variables. Then to recover the functions $g(t)$ and $g(x)$ back in the temporal or spatial domain from the corresponding function in the frequency domain, the *inverse Fourier transform* is used.

- In the time domain t:

$$g(t) = \int_{-\infty}^{\infty} G(f) e^{j2\pi ft} df \tag{2.13}$$

- In the spatial domain x:

$$g(x) = \int_{-\infty}^{\infty} G(k_x) e^{j2\pi k_x x} dk_x \tag{2.14}$$

Equations 2.11 and 2.13 represent the so-called *Fourier transform pair* in the time, which is often denoted as $g(t) \Leftrightarrow G(f)$. Likewise, Equations 2.12 and 2.14 is the *Fourier transform pair* in the spatial domain: $g(x) \Leftrightarrow G(k_x)$. The double arrow means that the right expression can be obtained by computing the forward Fourier transform of the expression in the left, and, similarly, the expression in the left can be obtained by taking the inverse Fourier transform of the expression in the right. Figure 2.23 shows some examples of Fourier transform pairs.

In the case of image processing, as images are finite (specific number of pixels in width and height), the Discrete Fourier Transform (DFT) will be used. This is the discrete version of the Fourier Transform (FT) that transforms a signal (or discrete sequence) from the time or spatial domain to its representation in the frequency domain. An important term to have in mind is the Fast Fourier Transform (FFT), which refers to any efficient algorithm that computes the DFT. Therefore, Equations 2.12 and 2.14 can be rewritten into its discrete form as follows:

- Discrete Fourier Transform (spatial domain):

$$G(k_x) = \sum_{x=0}^{N-1} g(x) e^{-j2\pi k_x x/N} \tag{2.15}$$

- Inverse Discrete Fourier Transform (spatial domain):

$$g(x) = \sum_{k_x=0}^{N-1} G(k_x) e^{j2\pi k_x x/N} \tag{2.16}$$

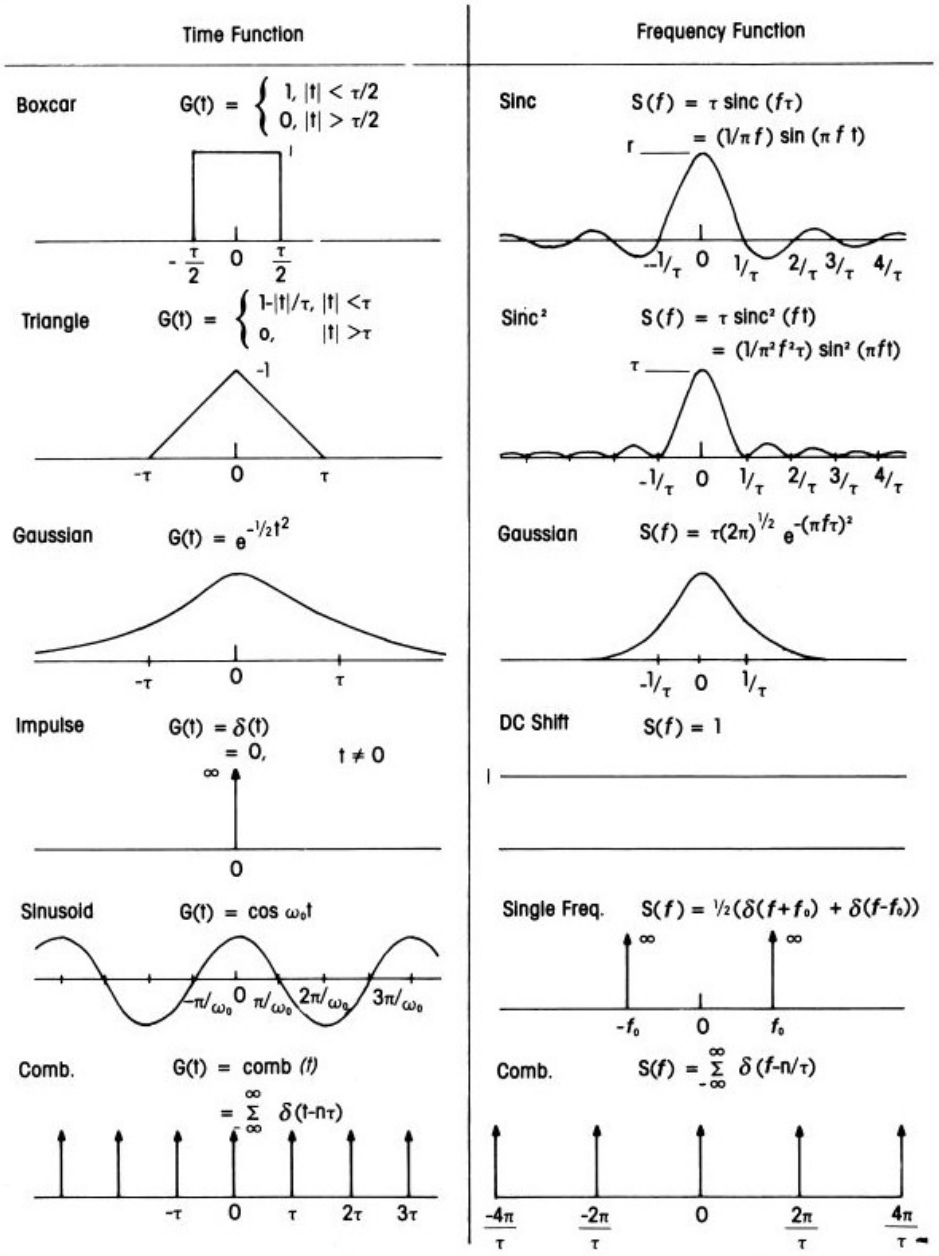

Figure 2.23: Examples of Fourier transform pairs. Source: [68].

The previous expressions can be re-written using Euler's formula $e^{jx} = cos(x) - jsin(x)$, where x is the term accompanying j in the exponent. If $g(x)$ is real (which in the case of an image, it will

be), the Fourier transform will generally have complex terms, with the form $G(k_x) = R(k_x)+jI(k_x)$, where $R(k_x)$ is the real part of the Fourier transform and $I(k_x)$ the imaginary part. The Fourier transform can also be expressed in the polar form as $G(k_x) = |G(k_x)|\, e^{j\varphi(k_x)}$, composed by a magnitude and phase for every frequency. The magnitude, phase angle and power of Fourier transform can be defined as follows.

- *Magnitude of the Fourier transform* $|G(k_x)|$, which is commonly called *Fourier spectrum* or *frequency spectrum* $|G(k_x)|$, is the square root of the sum of the squares of the real and imaginary parts.

$$|G(k_x)| = \sqrt{R(k_x)^2 + I(k_x)^2} \tag{2.17}$$

  The magnitude refers to the strength or amplitude of individual frequencies that contribute more to $g(x)$. It shows how much of a specific frequency is present in the signal. When visualized, the magnitude spectrum represents the distribution of amplitudes across various frequencies. If $g(x)$ is real, then the amplitude spectrum is symmetric around the Nyquist frequency (refer to Section 2.3.1.1 for further information about the Nyquist frequency). When plotting the magnitude $|G(k_x)|$, the x-axis represents the frequency and the y-axis is the amount of every frequency component in the signal. Rapidly changing functions require more high frequency content, whereas functions that are moving more slowly will have less high frequency components.
- *Phase angle of the Fourier transform* $\varphi(k_x)$. The phase spectrum denotes the shift or position of each frequency component with respect to a reference point, typically the origin. It indicates the timing or where each frequency's peak occurs within the signal. The phase information can significantly influence the signal's behavior, especially when reconstructing the original signal from its frequency components.

$$\varphi(k_x) = \arctan\left(\frac{R(k_x)}{I(k_x)}\right) \tag{2.18}$$

- *Power spectrum* $P(k_x)$. It is obtained as the squared magnitude and represents a metric of power spectral density (PSD), i.e. the power or energy associated with each frequency component.

$$P(k_x) = R(k_x)^2 + I(k_x)^2 \tag{2.19}$$

  The power spectrum allows for the identification of dominant frequencies or peaks in the signal, facilitating understanding the energy distribution across the frequency domain. Its application is particularly useful in multiple fields, such as signal processing, engineering and physics, where understanding the frequency characteristics and relative power of a signal is essential for analysis, filtering or modification.

Both the magnitude and phase provide a detailed representation of the frequency domain characteristics of a signal. They allow for the reconstruction of the original signal and analysis of its frequency content and temporal characteristics. The importance of each of these parts of a Fourier transform in image reconstruction will be further explained in Section 2.3.2.

Let's present an example to better understand how to relate the Fourier analysis to PIV image processing. Recall from Chapter 2.1 that, in PIV, due to diffraction-limited imaging, particle intensity can be approximated as a Gaussian distribution with standard deviation σ. Therefore, the corresponding Fourier transform pair is:

$$g(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{\left(-\frac{x^2}{2\sigma^2}\right)} \quad \Leftrightarrow \quad G(k_x) = e^{-2\pi^2\sigma^2 k_x^2} \tag{2.20}$$

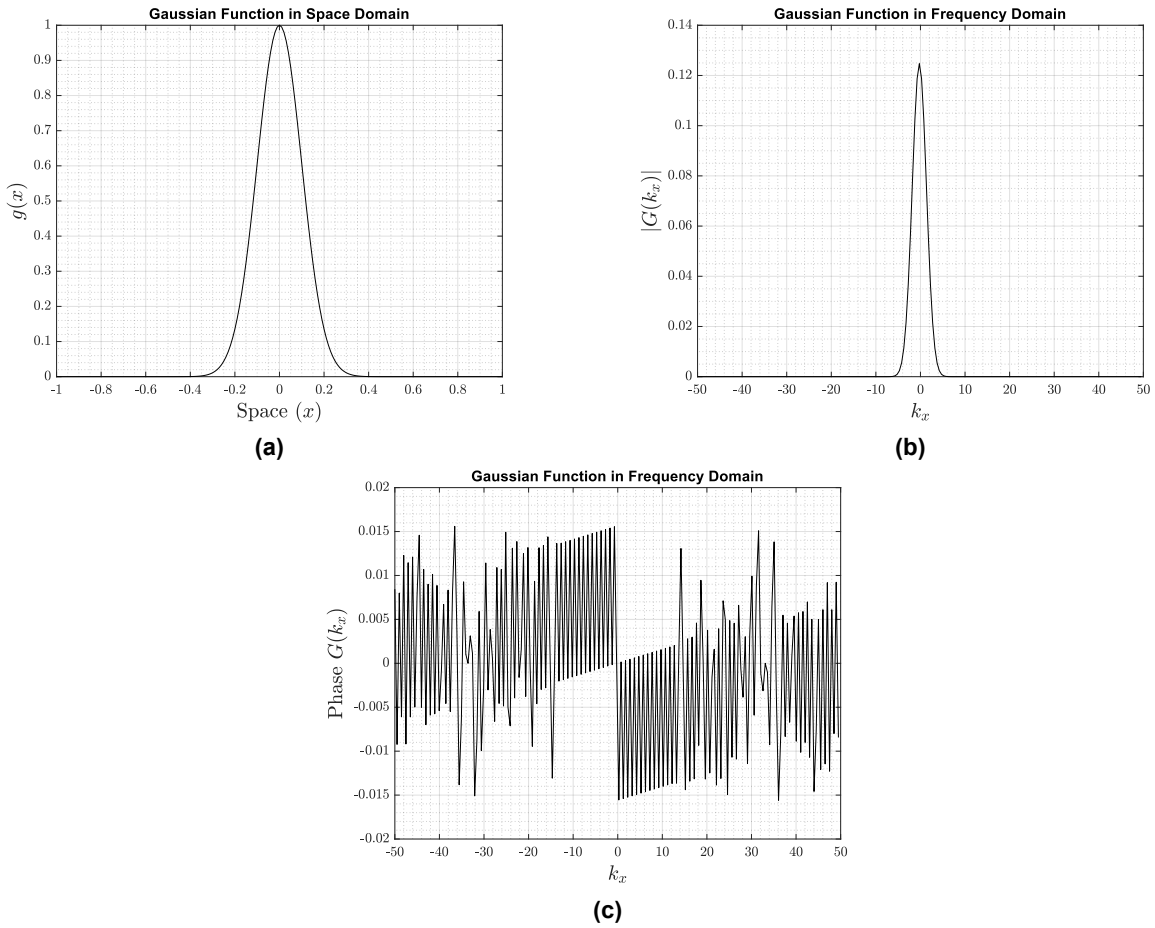

Figure 2.24: (a) Gaussian distribution with σ = 0.1, (b) its Fourier spectrum and (c) Fourier transform phase.

The Gaussian distribution $g(x)$, the magnitude of its Fourier transform (or Fourier spectrum) $|G(k_x)|$ and the phase are plotted in Figure 2.24. It can be observed that the Fourier transform of the Gaussian function is also a Gaussian (although missing the normalization constant). Moreover, notice that in Equation 2.20, the standard deviation σ moves from the denominator in the signal in the spatial domain to the numerator in the Fourier transform (frequency domain). This means that when the Gaussian distribution exhibits increased spread within the real spatial domain, its corresponding Gaussian Fourier transform in the frequency domain experiences a narrowing in its width, and conversely.

2.3.1.1. Nyquist frequency

The Nyquist frequency refers to the maximum frequency that can be accurately represented or sampled in a digitized signal. It is fundamental in the field of signal processing and is derived from the Nyquist-Shannon sampling theorem. The Nyquist theorem specifies that a sinuisoidal function in time or distance can be recovered with no loss of information as long as it is sampled at a frequency greater than or equal to twice the Nyquist frequency.

This means that for a discrete signal sampled at a frequency F_s, the highest frequency that can be represented without aliasing (where frequencies fold back incorrectly) is $F_s/2$, which is the Nyquist frequency. Frequencies above the Nyquist frequency would appear ”aliased” or

incorrectly reflected in the digitized signal, distorting the original frequency content. Therefore, to avoid aliasing and accurately represent frequencies in a digitized signal, the signal must be sampled at a rate at least twice the highest frequency component of interest to prevent information loss or distortion.

2.3.1.2. Convolution

Convolution of two functions entails flipping (rotating by 180°) one function with respect to its origin and sliding it past the other function. It is mathematically defined as the integral over all space of one function at τ multiplied another function at $x - \tau$, taken with respect to τ. The latter can represent any variable including time, frequency or even one, two or three dimensional space. Convolution will result in a function of a new variable τ, which will represent the same domain as the original variable x. This operation can be represented by a cross in a circle $\otimes$ (tensor product symbol) or by an asterisk $*$. In this document, the asterisk notation will be used to represent convolution. Given two continuous functions $g(x)$ and $h(x)$, with x being a continuous variable, the convolution is:

$$(g * h)(x) = \int_{-\infty}^{\infty} g(\tau)h(x - \tau)d\tau \tag{2.21}$$

In Equation 2.21, the flipping mentioned previously is given by the minus sign in $(x - \tau)$, where x is the displacement required to slide one function past the other, and τ is a dummy variable that is integrated out.

Then, considering the Fourier transforms $G(k_x)$ and $H(k_x)$ of $g(x)$ and $h(x)$, respectively, the Convolution Theorem [67] states that the Fourier Transform of the convolution of two functions, in this case in spatial domain, is the product of the Fourier Transforms of the functions. Conversely, if we have the product of the two transforms, we can obtain the convolution in the spatial domain by computing the inverse Fourier transform. In other words, $g * h$ and $G \cdot H$ are a Fourier transform pair and can be expressed as:

$$g(x) * h(x) \quad \Leftrightarrow \quad G(k_x) \cdot H(k_x) \tag{2.22}$$

This states that convolution in the spatial domain is analogous to multiplication in the frequency domain, the two being related by the forward and inverse Fourier transforms, respectively. This means that for linear, time-invariant systems, where the input/output relationship is described by a convolution, one can avoid convolution in the spatial domain by using Fourier Transforms. The convolution theorem is the foundation for filtering in the frequency domain: the real input signal can be modified by applying a filter $H(k_x)$ in the frequency domain and then apply inverse Fourier transform to recover the signal.

2.3.1.3. Filtering in frequency domain

A function in the temporal or spatial domain can be modified by filtering in the frequency domain. Given the Convolution Theorem described previously, this can be done by multiplication of a filter and the Fourier transform of the initial function. There are several types of filters used to modify frequency content and they can be generally classified into four main types (see Figure 2.25 for a visual example of each filter):

- *Low-pass filter*. This filter allows frequencies below a certain cutoff frequency to pass through and attenuates frequencies above that limit. It is typically used to remove high frequencies associated with noise or undesired components while keeping the lower-frequency components in a signal.
- *High-pass filter*. Opposite to a low-pass filter, a high-pass filter permits frequencies above a specific cutoff frequency and removes those below that cutoff. It is often used to eliminate low-frequency noise or to isolate high-frequency components.

- *Band-pass filter*. This type of filter allows a range or band of frequencies to pass through, while attenuating frequencies outside the specified range. It is employed when specific frequency bands within a signal are of interest, and the surrounding frequencies need to be filtered out.
- *Band-stop filter (low-high or notch filter)*. Opposite to a band-pass filter, this filter blocks a specific range or band of frequencies, allowing frequencies outside that range to stay in the signal. It is useful when particular frequencies need to be eliminated from a signal, but at the same time keeping the rest of the spectrum relatively unaltered.

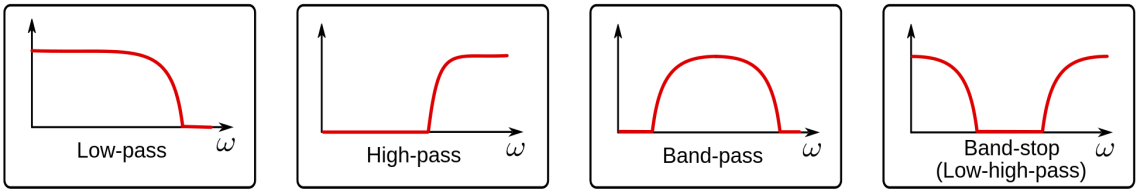

Figure 2.25: Types of filters. Source: [69].

Frequency domain filtering enables the selective manipulation (modify, enhance or isolate) of specific frequency components within signals or images, allowing for targeted enhancements, noise reduction, and various analysis techniques in diverse fields ranging from image and audio processing to telecommunications and beyond.

An illustrative example will be presented to enhance the understanding of how these filters work and their resulting outcomes. Low-pass and high-pass filters in the form of a Gaussian distribution with different values of standard deviation σ are applied to the previously presented Gaussian function distribution from Equation 2.20 depicted in Figure 2.24. Figures 2.26 and 2.27 show the low-pass and high-pass filters, respectively, with their resulting filtered Fourier transform and inverse Fourier transform signals.

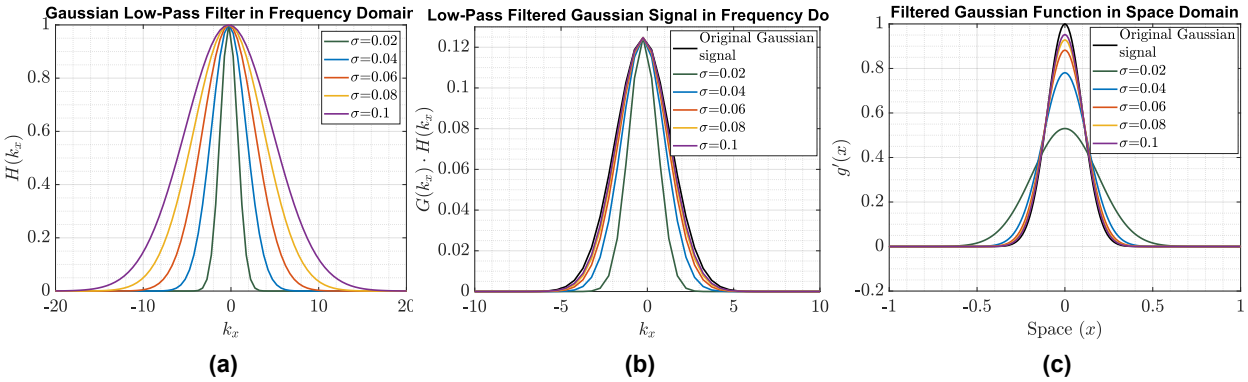

Figure 2.26: (a) Gaussian low-pass filter $H(k_x)$, (b) filtered Fourier transform $G(k_x) \cdot H(k_x)$ and (c) its inverse Fourier transform $g'(x)$ for different values of standard deviation σ.

In the case of a low-pass filter, as the standard deviation increases, the filter becomes more outspread, leading to more higher frequencies filtered out. Hence, a filter with high σ will suppress higher frequencies compared to a filter with lower σ. This is shown in the filtered Fourier transform signal in Figure 2.26 (b). As the filter size increases (higher standard deviation), the resultant filtered Fourier transform more closely resembles the original FT, which will result in the same initial function in space domain when the inverse FT operation is performed. This is depicted

in Figure 2.26 (c). A narrow filter results in a a filtered signal with decreased amplitude (-70% amplitude loss) compared to the original. Whereas, a more conservative low-pass filter (high standard deviation) will provide a signal that with more similarity to the input signal treated.

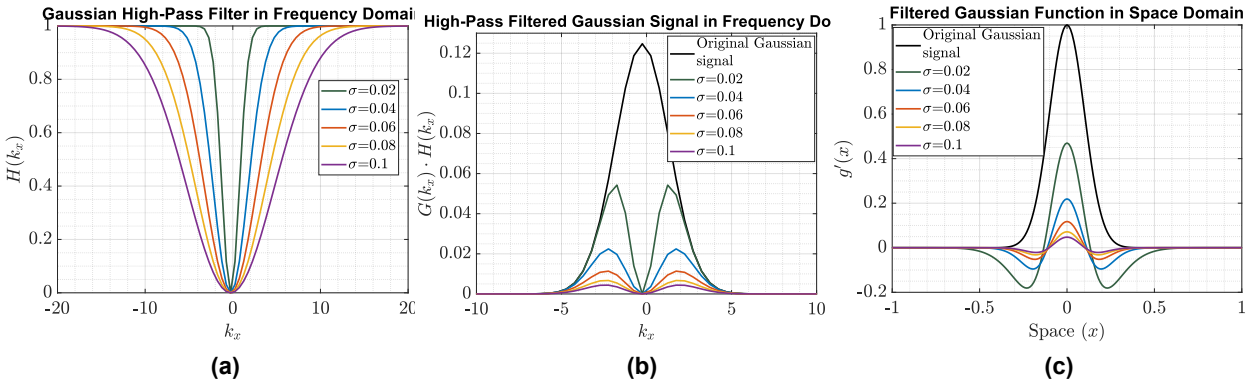

Figure 2.27: (a) Gaussian high-pass filter $H(k_x)$, (b) filtered Fourier transform $G(k_x) \cdot H(k_x)$ and (c) its inverse Fourier transform $g'(x)$ for different values of standard deviation σ.

Similar to the low-pass filter, as the standard deviation σ is increased, the high-pass filter H will filter out more high frequencies. Figure 2.27 (b) shows that as more frequencies are filtered out, the filtered signal amplitude decreases significantly until losing 96% of its amplitude with the highest standard deviation σ = 0.1. Even the narrowest filter (σ = 0.01) supposes a loss of 30% of the amplitude. This is because a high-pass filter is being applied, which means that the low frequency components that are related to the mean of the signal are being removed. A noticeable phenomenon that occurs with narrower high-pass filters is the negative overshoot that appears when the signal's amplitude approaches zero. This phenomenon is a consequence of the Gibbs effect, which will be further explained in the following Section 2.3.1.4.

2.3.1.4. Gibbs effect

In one-dimensional Fourier analysis, the Gibbs effect, or Gibbs phenomenon, refers to an overshoot or ringing artifact that occurs when approximating a discontinuous signal or a signal with sharp transitions using a finite number of Fourier components. This phenomenon happens because the Fourier basis functions are oscillatory and, hence they cannot represent or contain sharp transitions or discontinuities. However, when these are present in the modified Fourier transform and the inverse operation of the Fourier transform is performed, the resulting signal overshoots near the discontinuity, leading to oscillations that do not converge to the true value. This is precisely what happened in the previous example in Figure 2.27. The filtered or modified Fourier transforms with a low standard deviation high-pass filters present a discontinuity near the zero frequency. When the inverse Fourier transform is obtained the resulting signals overshoot near the discontinuity.

Another example would be when representing a signal with sharp changes, such as a square wave, using a Fourier transform with a finite number of terms. The reconstructed signal might manifest oscillations near the discontinuities as depicted in Figure 2.28.

The Gibbs phenomenon is a fundamental limitation when approximating discontinuous signals using a finite number of Fourier components and is an inherent property of Fourier analysis when dealing with signals that have sharp transitions or discontinuities. Albeit the oscillations near the discontinuities cannot be completely eliminated, the amplitude of these can be reduced by using a higher number of Fourier components in the approximation.

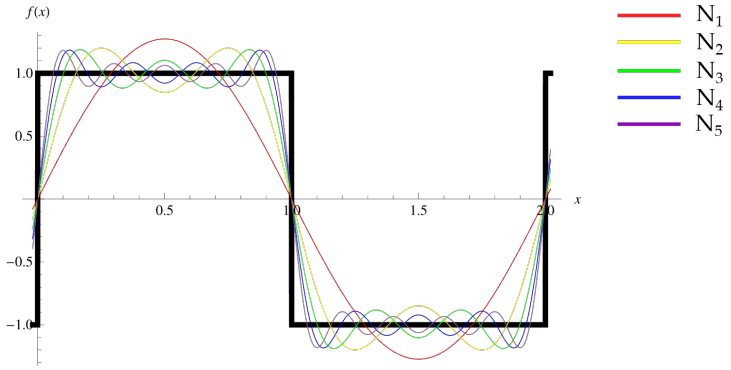

Figure 2.28: Gibbs effect example on a rectangle pulse (where N is the number of terms used to approximate the pulse, with $N_5 > N_4 > N_3 > N_2 > N_1$). Source: [70].

2.3.2. 2D Fourier transform

The 2D Fourier transform is an extension of the 1D Fourier transform, used to analyze two-dimensional signals, such as images. It is a significantly important tool in image processing, used to decompose an image into its sine and cosine components. The result represents the image in the Fourier or frequency domain, contrasting with the original input image, which exists in the spatial domain.

- Fourier Transform:

$$G\left(k_x, k_y\right) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) e^{-j2\pi(k_x x + k_y y)} dx dy \tag{2.23}$$

- Inverse Fourier Transform:

$$g(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G\left(k_x, k_y\right) e^{j2\pi(k_x x + k_y y)} dk_x dk_y \tag{2.24}$$

where k_x and k_y are the frequencies associated to the x (horizontal) and y (vertical) axis, respectively.

As mentioned in the previous Section, the Discrete Fourier Transform (DFT) is used to analyze discrete signals, such as digital audio and images. It decomposes a discrete signal into its individual frequency components, allowing for the analysis of the frequency content of the signal. The DFT and its inverse operation in the 2-dimensional space are computed as follows:

- Discrete Fourier Transform:

$$G\left(k_x, k_y\right) = \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} g(x, y) e^{-j2\pi\left(\frac{k_x}{N} x + \frac{k_y}{M} y\right)} \tag{2.25}$$

  where $g(x, y)$ is an image of size (NxM), k_x the spatial frequency in the x-axis and k_x the spatial frequency in the y-axis. These equations are evaluated for values of the $k_x = 0, 1, 2, ..., N-1$ and $k_y = 0, 1, 2, ..., M-1$.

- Inverse Discrete Fourier Transform:

$$g(x, y) = \frac{1}{NM} \sum_{k_x=0}^{N-1} \sum_{k_y=0}^{M-1} G\left(k_x, k_y\right) e^{j2\pi\left(\frac{k_x}{N} x + \frac{k_y}{M} y\right)} \tag{2.26}$$

  for $x = 0, 1, 2, ..., N-1$ and $y = 0, 1, 2, ..., M-1$.

The DFT is the sampled Fourier Transform and therefore does not contain all frequencies forming an image, but only a set of samples which is large enough to fully describe the spatial domain image. The number of frequencies corresponds to the number of pixels in the spatial domain image, i.e. the image in the spatial and Fourier domain are of the same size.

2.3.2.1. Properties of the 2D Fourier transform

Linearity

The principle of superposition states that the response produced by the combined effect of several inputs on a system equals the sum of the individual responses that each input would cause independently. Hence, a system is linear if its response to two signals is equal to the sum of the responses of the individual signals. A system that satisfies the principle of superposition is linear. In the case of a Fourier transform, it is a linear operation since one can affirm that:

$$\mathcal{F}(Ag_1 + Bg_2) = A\mathcal{F}(g_1) + B\mathcal{F}(g_2) \tag{2.27}$$

where $\mathcal{F}$ denotes the Fourier transform operation and g_1 and g_2 represent two signals. This equation shows that if you take the Fourier transform of a linear combination or sum of functions, the resulting Fourier transform is equal to the sum of their individual Fourier transforms.

The application of the linearity property in the Fourier transform allows analysis and manipulation of signals in the frequency domain. It enables the decomposition of complex signals into simpler sinusoidal components in the frequency domain, enabling a deeper understanding of the frequency content of the signals for its analysis and reconstruction. This means that images can be divided by looking at their frequency domain content and identifying the corresponding components.

Shift invariance Shift invariance, also known as translation invariance, in Fourier analysis refers to a property where a system or an operation remains unaffected by shifts or translations in the input signal's domain, specifically concerning time or space. In the context of Fourier analysis, a system or operation is considered shift-invariant when the Fourier transform of a translated or shifted input signal is directly related to the Fourier transform of the original signal, with the same shift applied to the transformed signal.

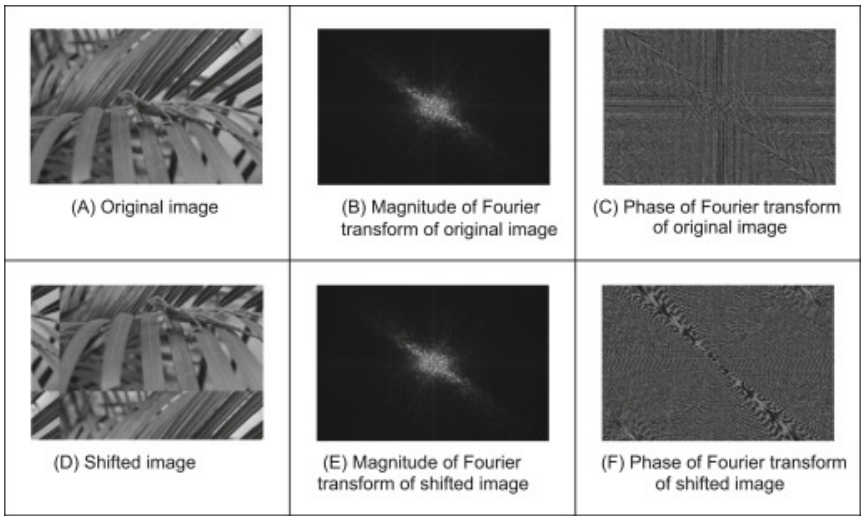

Figure 2.29: Example of shift invariance in Fourier transform. Source: [71].

This property essentially means that if all features in an image are shifted or if this image is captured from a different position, the magnitude of the Fourier transform will remain unchanged, but there will be a phase shift in the frequency domain. Shift invariance is a significant property because it ensures that the phase change caused by a shift is consistent across all frequencies.

Shift invariance is vital in many signal processing and analysis applications. For instance, in image processing, a shift-invariant system ensures that features or patterns in an image are identifiable and can be recognized even if the image is shifted. This property enables various algorithms, filters, and operations to be more robust and effective in analyzing signals or images regardless of their spatial or temporal position.

Low and high frequency components In Fourier analysis, the frequency components of a signal can be categorized into low and high frequencies, each carrying distinct characteristics.

Low frequencies are associated with slower oscillations or gradual changes in a signal. In the frequency domain, they are near the origin as shown in Figures 2.30 (a) and (b). The origin is commonly known as the *dc component* and it is where the average value of the original image is contained. For instance, in an image, low-frequency components represent smoother transitions between intensities. Lower frequency components contribute to the base or background of the signal.

On the other side, high frequencies are situated further from the origin and imply rapid changes in pixel intensity as can be sharp edges or features in the image (see Figures 2.30 (c) and (d)). These pertain to faster oscillations or rapid changes in a signal. High frequencies in the frequency domain represent the components of a signal that change rapidly or have shorter cycles. They capture fine details or sharp changes in images, such as edges or textures.

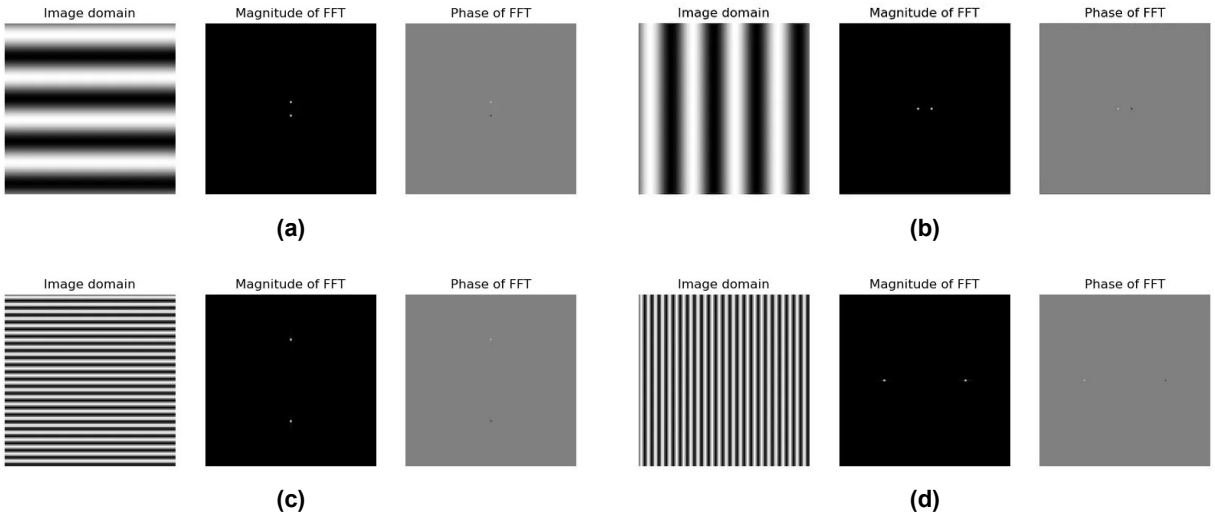

Figure 2.30: Examples of sinusoid gratings with different frequency and direction. Source: [72].

Note in the figure above that the Fourier transform components appear in the opposite direction of their corresponding original image. This is because the change in pixel intensity occurs in the contrary direction compared to the horizontal or vertical bars. For instance, Figure 2.30 shows horizontal lines that change intensity in the vertical direction, thus the Fourier transform shows the frequency components in the vertical axis.

Symmetry Another property of a Fourier transform is that if a function $g(x, y)$ in the spatial domain is real, its Fourier transform is a conjugate symmetric, thus

$$G^*(k_x, k_y) = G(-k_x, -k_y) \tag{2.28}$$

This can also be applied on the opposite case: when the function $g(x, y)$ in the spatial domain is imaginary, its Fourier transform is a conjugate antisymmetric: $G^*(-k_x, -k_y) = -G(k_x, k_y)$. This will be important in image processing when filtering in the frequency domain and then reconstructing the filtered function in the spatial domain via the inverse Fourier transform operation. An image is generally real *per se*, hence if a filter is applied in its Fourier transform, the output of the inverse FT will have to be real as well. In consequence, the filtered FT must be a conjugate symmetric for the resulting image to be real.

In Matlab, the function *ifft2()* is used to perform the inverse Fourier transform in a 2D variable (i.e. an image). Its input is the filtered Fourier transform and, optionally, if the *'symmetric'* flag is used as second argument, it will treat the filtered FT as conjugate symmetric by ignoring the second half of its elements (that are in the negative frequency spectrum). This option is useful when the filtered Fourier transform is not exactly conjugate symmetric, merely because of round-off error. If this option is not used, the resulting image would contain imaginary elements, which would not represent properly the filtered image.

Fourier transform components Similar to the 1D case, the 2D Fourier transform is generally complex and can be expressed as

$$G(k_x, k_y) = R(k_x, k_y) + jI(k_x, k_y) = |G(k_x, k_y)|\, e^{j\varphi(k_x, k_y)} \tag{2.29}$$

(where R and I are the real and imaginary parts of $G(k_x, k_y)$. Therefore, in order to represent its components, the magnitude or Fourier spectrum $|G(k_x, k_y)|$ and the phase angle or phase spectrum $\varphi(k_x, k_y)$ are visualized. Their calculation is analogous to the process shown in Section 2.3.1 in Equations 2.17 and 2.18. However, as a reminder and for clarification, the corresponding equation in 2D space are presented below.

$$|G(k_x, k_y)| = \sqrt{R(k_x, k_y)^2 + I(k_x, k_y)^2} \qquad \varphi(k_x, k_y) = \arctan\left[\frac{R(k_x, k_y)}{I(k_x, k_y)}\right] \tag{2.30}$$

From Equation 2.25, and considering the origin of the frequency domain ($k_x = 0$, $k_y = 0$), the discrete Fourier transform is:

$$G(0, 0) = \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} g(x, y)$$

This proofs that the zero-frequency component DFT is proportional to the average of the signal in the spatial domain g(x, y), as

$$G(0, 0) = NM \frac{1}{NM} \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} g(x, y) = MN\overline{g}$$

where $\overline{g}$ represents the average value of a function $g(x, y)$. Additionally,

$$|G(0, 0)| = MN\,|\overline{g}| \tag{2.31}$$

As the proportionality constant MN is usually large, $|G(0, 0)|$ typically is the largest component of the spectrum by several orders of magnitude larger than the rest of the content. Since the

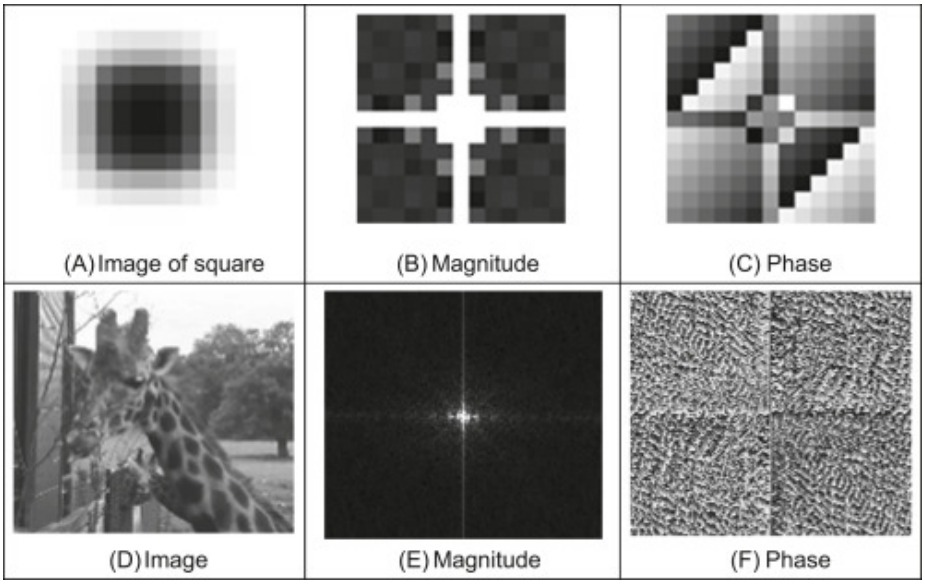

Figure 2.31: Two examples and their Fourier and phase spectrums (from left to right). Source: [71].

frequency components k_x and k_y are zero at the origin, the term $|G(0,0)|$ is commonly called the *dc component* of the transform.

All these concepts might seem somewhat abstract, hence, Figure 2.31 serves as an illustration for a visual interpretation of what magnitude and phase represent in the image Fourier domain.

The magnitude shows how much signal there is at a particular frequency component and the phase encodes the spatial information (indirectly) about how the image features are distributed. In the magnitude plot in Figure 2.31 (B) and (E), the area around the origin of the magnitude contains the highest values (and thus appears brighter in the image). As explained previously, this area is the *dc* component and includes the average of the input image. Because this term dominates, the dynamic range of other intensities in the displayed image appear rather compressed (shown clearer in (E)).

But, which term is more relevant when reconstructing an image: the magnitude or the phase? This question can be answered with the example in Figure 2.32, which shows two examples of input images, their corresponding Fourier and phase spectrums and, additionally, reconstruction of the images from only using the magnitude and only the phase.

The spectral components of the Discrete Fourier Transform (DFT) dictate the amplitudes of the sinusoidal elements that constitute an image. A higher amplitude at a particular frequency indicates a greater influence of the corresponding sinusoid within the image, while a lower amplitude suggests its reduced presence. If the inverse DFT of the image example presented before is done only based on the Fourier spectrum (setting the phase angle to 0), the resulting images (Figure 2.32 (b)) exclusively retain data about the intensity of the pixels, with the *dc* term being the most dominant. However, the resulting image lacks shape details due to the phase being set to zero; the pixels are not arranged with a coherent form, not giving any information to the user for self-interpretation of the image.

Although less straightforward, the phase components hold similar significance as the magnitude. They denote the positional displacement of the sinusoidal constituents from their origins. Although there is no detail in the phase that would lead us by visual analysis to associate it with the structure of its corresponding image, the information they provide to image reconstruction is crucial in determining shape features of the image. To illustrate this, the input images in Figure

2.32 were reconstructed using only its phase angle (computing the inverse DFT using $\varphi(k_x, k_y)$, but setting $|G(k_x, k_y)| = 1$). It is can be seen that much of the intensity information has been lost since the information is carried by the spectrum since it was not used in this reconstruction. However, the shape features from the original raw image are clearly maintained and the faces of the original images can be clearly recognised, even though with no such intensity.

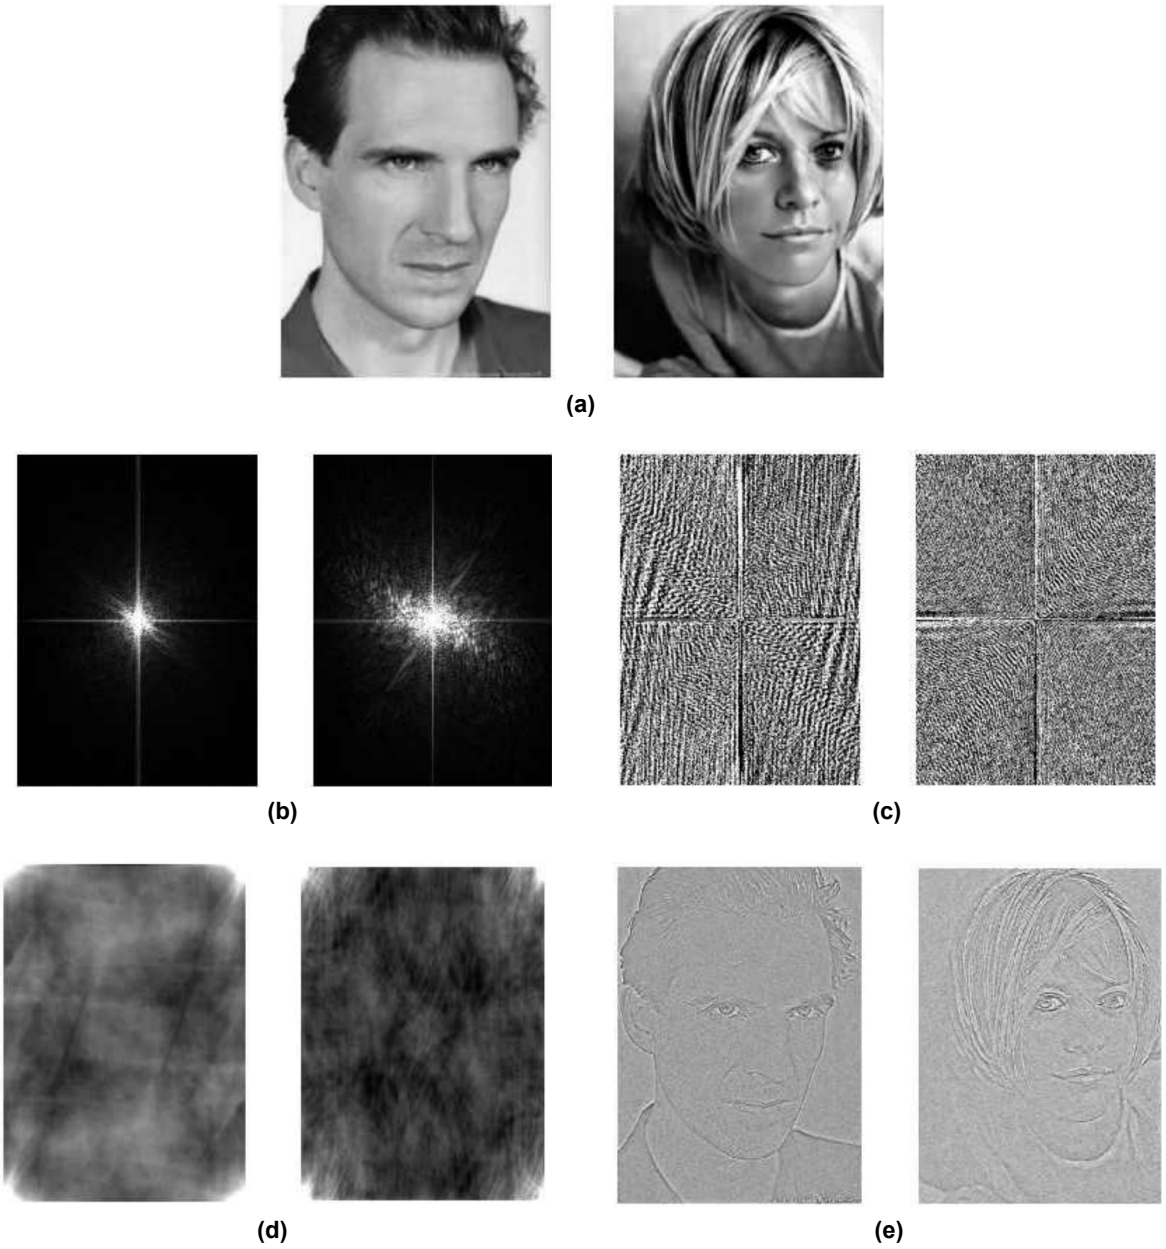

Figure 2.32: (a) Input images, (b) Fourier spectrum $|G(k_x)|$, (c) phase angle, (d) images reconstructed using only their spectrum and (e) images reconstructed using only their phase angle. Source: [73].

As a further example, Figure 2.33 show two examples of what occurs if an image is reconstructed when the magnitude is maintained but the phase is swapped with the phase information of another image. The left images show the original images of a cameraman and a girl. In the right hand side, there are the reconstructed images. The top one is regenerated with the cameraman magnitude and the phase of the girl's image. The features shown are clearly from the girl's image; no information about the cameraman can be interpreted in this image. Similarly, the bottom one

is again regenerated, but with the girl's magnitude and the cameraman's phase. Again, only the shape of the camera man can be distinguished. This strongly illustrates the importance of the phase angle in determining shape characteristics in an image.

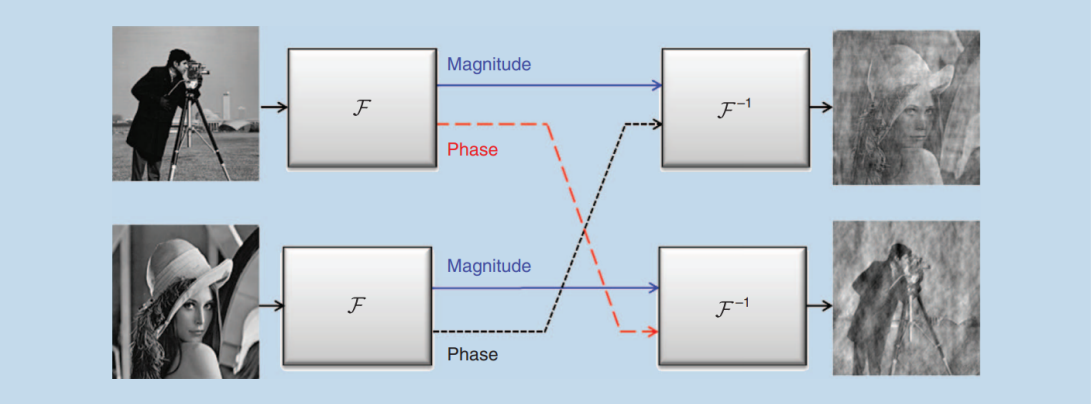

Figure 2.33: Two images are Fourier transformed. Their phases are swapped and then, they are inverse Fourier transformed. Source: [74].

To sum up, the magnitude of the 2-D DFT contains the intensities or the amount of a specific frequencies within the image, while the associated phase represents the angles that provide positional information about the spatial localization of features within the image. Note from the previous examples that, in general, visual interpretation of phase angle images yields little intuitive information. However, it has been shown that the phase is extremely important when reconstructing an image as it provides powerful information of the features and where they are located within an image.

2.3.2.2. Nyquist frequency

Analogously to 1D (see Section 2.3.1.1), the Nyquist frequency in 2D space is the maximum spatial frequency that can be accurately represented or sampled in an image in both the horizontal (x-axis) and vertical (y-axis) directions. It is significantly important as it ensures that the image is properly sampled without aliasing, where higher frequencies fold back or create artifacts due to undersampling.

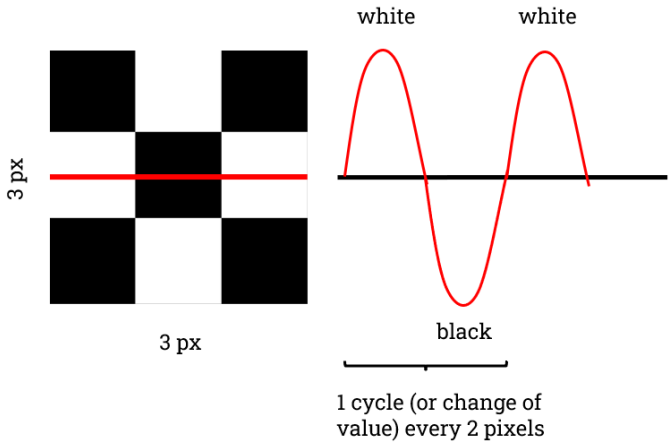

Figure 2.34: Nyquist frequency representation in 2D.

For a 2D digital image, the Nyquist frequency is related to the maximum frequency content that can be accurately captured during the digitization process. For this, if an image is sampled at a rate of f_s samples per second (or per unit distance) in the x and y directions, the Nyquist frequency will be $f_s/2$ cycles per unit distance in both horizontal and vertical dimensions. When the function domain is distance, as in the case of image processing, the sample rate might be pixels/cm and the corresponding Nyquist frequency would be in cycles/cm.

2.3.2.3. Filtering in 2D spatial frequency domain

To explain filtering, recall convolution from Section 2.3.1.2 applied to 1D signals. Extrapolating the expression to 2D results in the called *2D circular convolution*:

$$(g * h)(x, y) = \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} g(n, m)h(x - n, y - m) \tag{2.32}$$

for $x = 0, 1, 2, ..., N - 1$ and $y = 0, 1, 2, ..., M - 1$. The 2D Convolution theorem states that

$$g(x, y) * h(x, y) \quad \Leftrightarrow \quad G(k_x, k_y) \cdot H(k_x, k_y) \tag{2.33}$$

$$g(x, y) \cdot h(x, y) \quad \Leftrightarrow \quad \frac{1}{NM} G(k_x, k_y) * H(k_x, k_y) \tag{2.34}$$

where $G(k_x, k_y)$ and $H(k_x, k_y)$ are the the Fourier transforms of the 2D space functions $g(x, y)$ and $h(x, y)$. Equation 2.33 represents an equivalence between the spatial and frequency domain and is considered the basis of linear filtering in the frequency domain.

Filtering in the frequency domain consists of modifying the Fourier transform of an image by multiplying the latter by a filter or filter function $H(k_x, k_y)$, then computing the inverse transform to obtain the spatial domain representation of the processed result. Frequency is directly related to spatial rates of change, it is not difficult intuitively to associate frequencies in the Fourier transform with patterns of intensity variations in an image. As proofed in Section 2.3.2.1, the slowest varying frequency component $(k_x, k_y) = (0, 0)$ is proportional to the average intensity of an image. As we move away from the origin of the transform, the low frequencies correspond to the slowly varying intensity components of an image. As we move further away from the origin, the higher frequencies begin to correspond to faster and faster intensity changes in the image. These are the edges of objects and other components of an image characterized by abrupt changes in intensity.

Filtering techniques in the frequency domain are based on modifying the Fourier transform to achieve a specific objective, and then computing the inverse DFT to return to the spatial domain. The two components of the transform that can be interpreted are the transform magnitude (spectrum) and the phase angle. Visual analysis of the phase component generally is not very useful, whereas the spectrum provides some useful guidelines as to the gross intensity characteristics of the image from which the spectrum was generated.

2.3.2.4. Applications of image filtering in frequency domain

Image processing using the Fourier transform in the spatial domain finds applications in various fields due to its ability to analyze, enhance, and manipulate signals and images by understanding their spatial frequency content. It enables the analysis of an image's frequency content, revealing details and patterns that might not be easily discernible in the spatial domain. Some significant fields where this technique is applied include:

- Medical Imaging. Fourier-based image processing is used in medical imaging techniques like MRI, CT scans, and ultrasound to enhance image quality, reduce noise, and improve

diagnostic capabilities by isolating and analyzing specific features in the images ([75], [76], [77]).

- Artificial Intelligence and Computer Vision. Fourier-based image processing is used in computer vision tasks, like object detection, segmentation, and feature extraction, contributing to machine learning and AI algorithms for various applications. It also helps in biometric image enhancement and feature extraction for facial or fingerprint recognition systems [78], [79], [80].
- Image Compression and Storage. Fourier techniques are employed in image compression methods (e.g., JPEG) to reduce file size while retaining essential information, making it feasible to store and transmit images efficiently ([81], [82]).
- Astronomy and Astrophysics. In astronomy, Fourier-based image processing assists in cleaning and enhancing astronomical images, enabling clearer observations and analysis of celestial objects ([83], [84]).

Image processing via Fourier Transform in the spatial domain is a powerful tool that finds applications across multiple fields, contributing to better analysis, enhancement, and interpretation of images for various purposes, including medical diagnostics, security, artificial intelligence, and more.

3

Characterization of reflections

Recall that the aim of the project is to reduce or mitigate the impact of reflections while preserving the pixel intensity of particles in order to improve particle tracking in robotic PIV. To achieve this, it is crucial to understand what are the main characteristics associated with reflections and those linked to particles. Characterizing laser light reflections involves analysing how they appear in the images acquired by the imaging system and in the resulting processed Shake-the-box data. Therefore, it is essential to perform a reflection characterization prior to proposing a methodology.

3.1. Reflection characterization in PIV images

The shape of the reflection highly depends on the laser beam characteristics. Figure 3.1 shows a set of examples of raw images with reflections for two different PIV techniques.

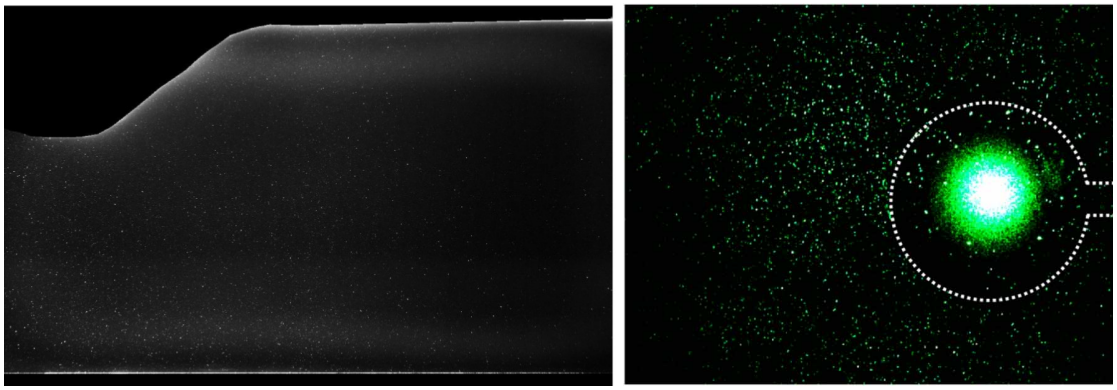

(a) **(b)**

Figure 3.1: Laser light reflection on (a) a planar PIV case, (b) robotic PIV case. Source: [85], [86].

In the case of 2D planar PIV (Figure 3.1 (a)), the laser beam is more concentrated and, therefore the resulting reflections appear as concentrated areas or sharp edges. Whereas in the case of 3D robotic PIV, as a volume is to be acquired, the laser beam is expanded in a conical shape, which results in a large, diffused region of high intensity as depicted in Figure 3.1 (b). The reflection can easily be spotted, as a glare spot, of high variable intensity, and much bigger than the particles, which, due to diffraction effects, typically appear in a Gaussian shape with a small diameter of a few pixels. In the worst case, it can saturate the detector but in general,

the reflection intensity is higher or similar than that of the particles, which will have a significant impact on the resulting velocity field. This will be important when trying to distinguish them from particle images.

Let's take a closer look at the sphere case of Figure 3.1 (b). A close-up on the reflection and particle images is shown in Figure 3.2.

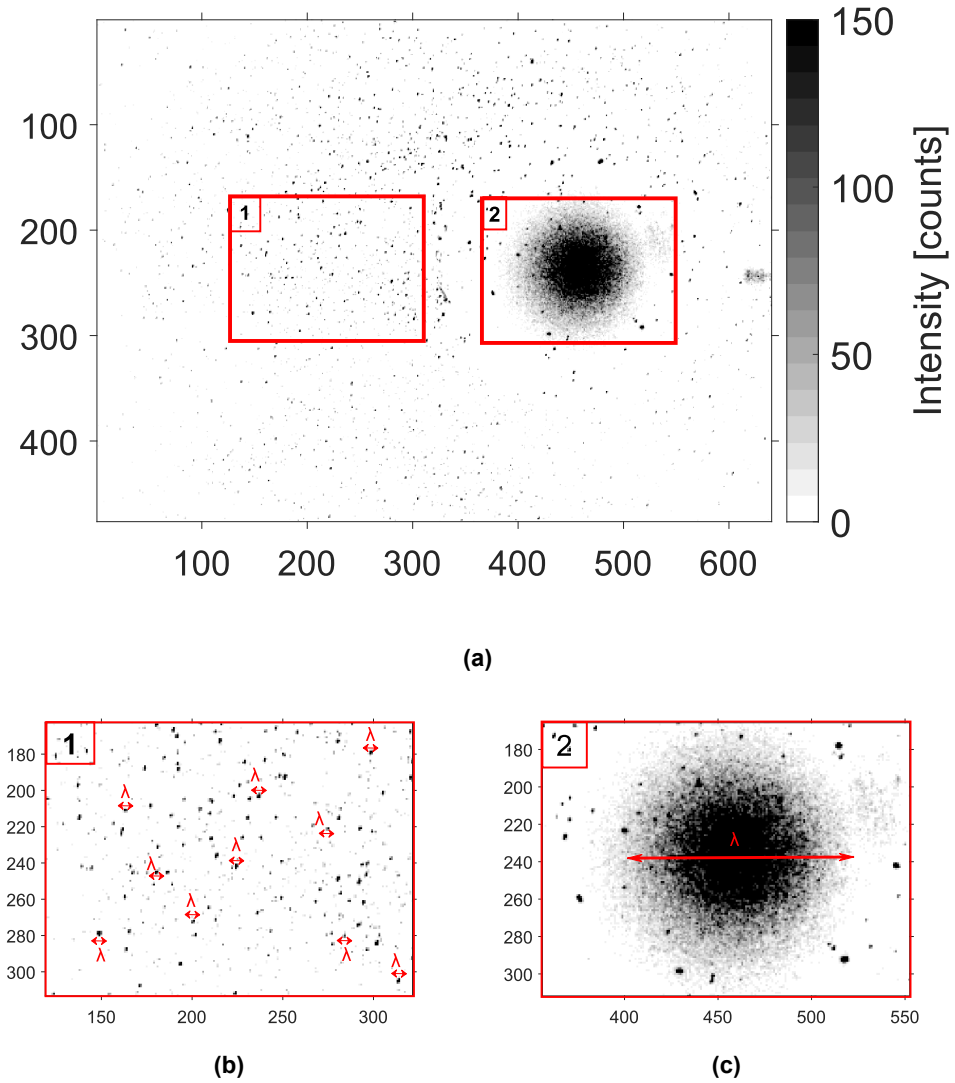

Figure 3.2: (a) Instantaneous Robotic PIV raw image of sphere case, close-up on a region (b) with particles and (c) the reflection.

The main and, in general, most obvious characteristic of reflections is their size: they appear as large areas of high pixel intensity, often to the point of saturation of the camera detector. This characteristic makes them visually prominent in images. On the other side, particles appear as small areas of high pixel intensity (around a few pixels) that resembles a Gaussian distribution. Filtering by pixel intensity is not accurate since both particle images and reflection have similar intensity values, but there is a clear difference in the size that allows us to distinguish reflections and particles by their respective wavelengths. Reflections tend to have longer wavelengths, while particles, being smaller, have much shorter characteristic wavelengths, as shown in Figure 3.2.

3.2. Reflection characterization in Shake-the-Box data

From the literature review, it has been shown that occasionally the image pre-processing techniques fail to fulfil the task of fully removing reflections. When this occurs, the resulting processed 3D data appears highly affected by their presence, leading to inconsistencies and potentially affecting negatively its interpretation. To understand how reflections appear in the processed Shake-the-Box data, two cases are compared in Figure 3.3: processed data from raw images (no image pre-processing applied) and from pre-processed images with the Butterworth time filter (considering a filter length of 9 images).

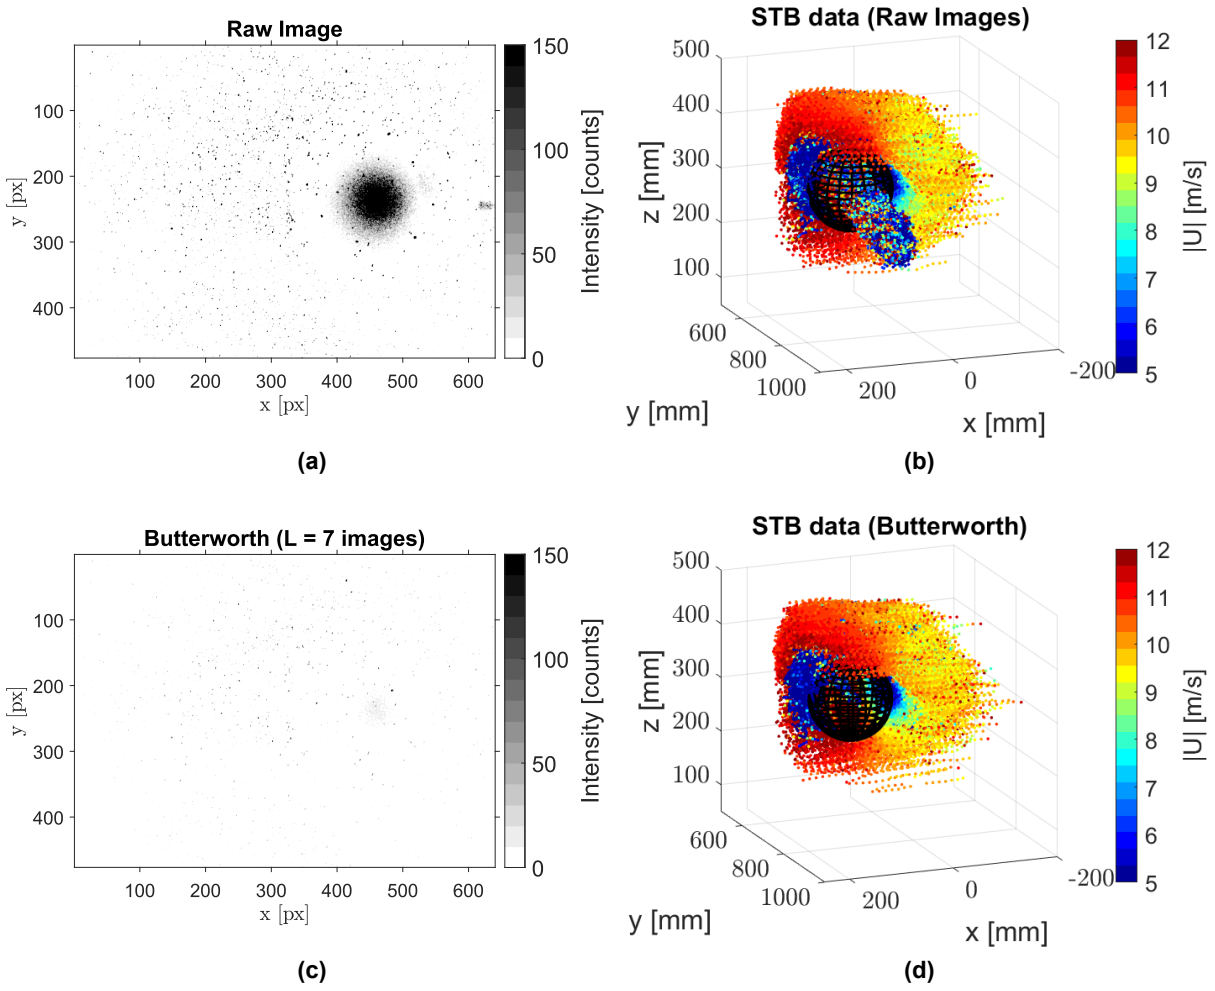

Figure 3.3: Left: Raw and time-filtered (Butterworth) images (top to bottom). Right: corresponding Shake-the-Box data.

The raw image case in Figure 3.3 (a-b) shows that the reflection appears as an accumulation of particles in a cylinder shape that propagates along the line-of-sight of the camera. These particles are often referred to as ghost particles or false positives, because as the reflection contains pixels with similar intensity as particles, the particle tracking algorithm confuses these as particles inside the reflection region, although no particles are actually present. This misinterpretation is clearly proven by looking at the Butterworth case, where the reflection is accurately fully mitigated (thanks to its steadiness over time). The processed STB data shows no or little evidence of the presence of a reflection and, as a result, allows to track particle tracers that cross over the reflection region.

These particles are completely lost in the raw image processed case. Therefore, it is essential to mitigate reflections, as they yield misleading and confusing results. These artifacts can distort data interpretation, potentially leading to erroneous conclusions.

By human eye, it is fairly simple to identify the location of the reflection by its shape and distinct behaviour compared to its environment. But, is there a variable that can be used to distinguish regions of reflections from real particle tracks? Particles detected inside reflection regions can appear with a wide spectrum of velocities: either high or low. Hence, analysis of the velocity information fails to yield definitive distinctions between particles and reflections. Consequently, an alternative parameter should be investigated as potential reflection-particles distinguishable characteristic in 3D data.

Recall from Section 2.1.1.2.1, binning is used as an averaging processing of the whole set of data. This binning averages the track information into bins, which results in a set of parameters that can be checked.

- Acceleration
- Standard deviation
- Number of particles
- Reynolds stresses
- Average kinetic energy
- Turbulent kinetic energy
- Turbulent shear stress
- Uncertainty for mean values

All of them are closely related to the motion of the tracks, except for one of them: the number of particles. The latter is accounted by the tracer particles Shake-the-Box is able track over time. In each bin, this value is averaged for all time steps and the particle concentration can be computed (given the bin size and the number of images processed). Figure 3.4 shows the particle concentration on the raw image and the Butterworth cases studied previously, considering the following binning parameters: 40x40x40 voxels with 0% overlap and second order polynomial approximation.

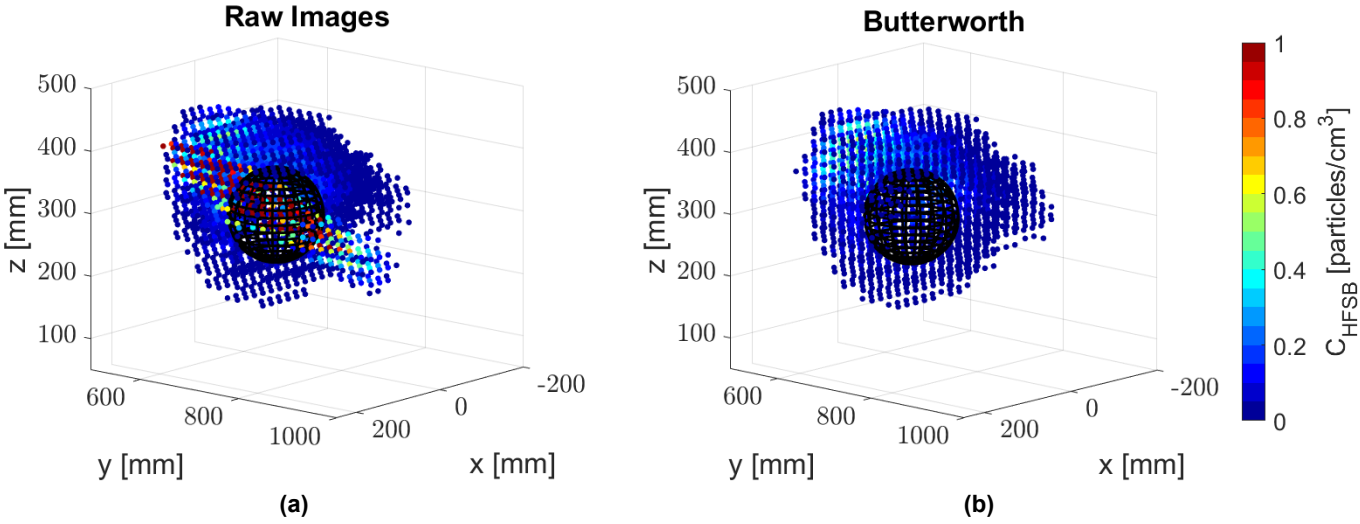

Figure 3.4: Particle concentration on (a) raw image case and (b) time-filtered (Butterworth) case.

The binning performed on the processed raw images shows that the particle concentration on the reflection has distinguishable higher values compared to the rest of the domain. This indicates that the particle concentration parameter could be potentially used as variable that allows to distinguish between reflection and real particles within the flow field.

4

Reflection treatment techniques

This chapter introduces three novel approaches developed with the aim to improve the reliability and accuracy of PIV data by addressing light reflection in both PIV images and 3D data. The data from an experimental database of a flow around a sphere (for further details, refer to [86]) is employed to verify the validity of the principles. The pre-processed images from a single view acquired with the Robotic PIV system are analysed with the Shake-the-Box particle tracking algorithm.

4.1. Reflection attenuation via Spatial Fourier Filtering

Section 2.3 presented an overview of the diverse applications of Fourier Analysis in various fields. Now, let's explore its potential application in the PIV (Particle Image Velocimetry) image processing pipeline. As described in Section 2.3.2, spatial Fourier frequency analysis can be employed as an image pre-processing technique. This allows to transform an image in real space into its distribution of spatial frequencies. This spatial frequency can also be called *wavenumber* and is represented with the Greek letter κ.

$$\kappa = \frac{1}{\lambda} \tag{4.1}$$

Equation 4.1 defines the wavenumber the inverse of the wavelength $[px^{-1}]$ of a given signal. Images in wavenumber space are obtained by means of the Fourier transform operation. Reflections are big in general, thus will tend to have a much larger wavelengths compared to particle images, which normally appear with a particle size/diameter of approximately 3-5 pixels.

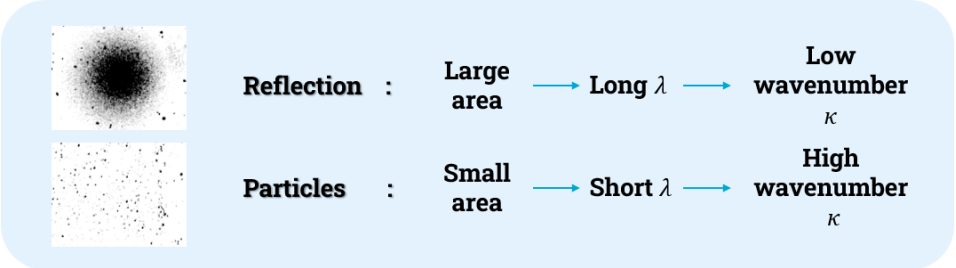

Figure 4.1: Reflection-particle wavelength comparison.

Consequently, particles will contain high wavenumber information, whereas reflections will be represented by lower wavenumbers. Hence, in order to filter out reflections and keep particle content, a high-pass filter must be used on the Fourier space. Let's visualize this effect with an example in 1D. Three different cases are displayed: a unit impulse signal, a Gaussian with low standard deviation (representing a particle) and a Gaussian with high standard deviation (representing a reflection).

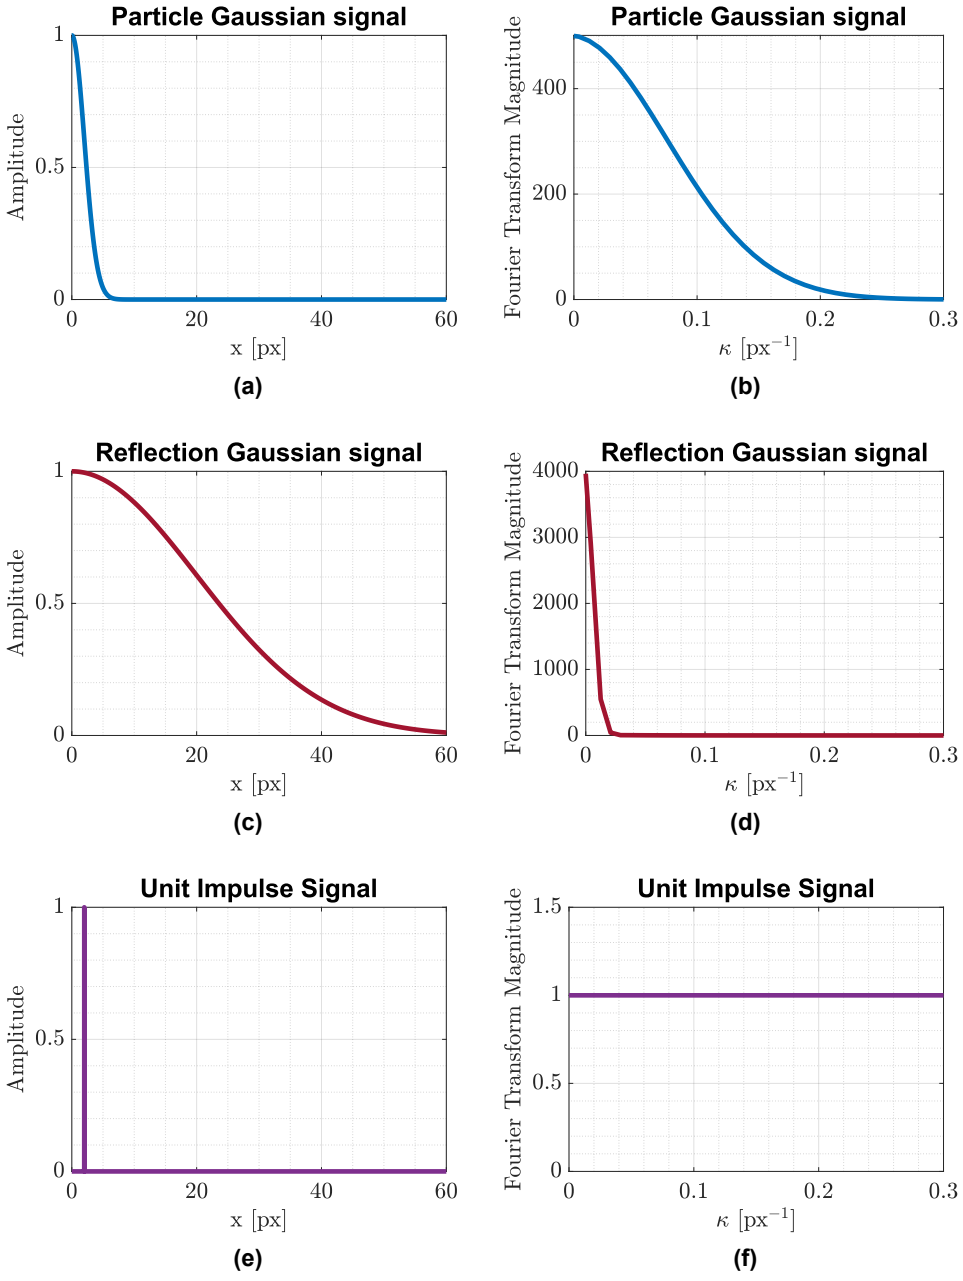

Figure 4.2: Left: particle, reflection and unit impulse signals (top to bottom). Right: The corresponding Fourier transforms.

Figure 4.2 (a) and (c) depict a narrow (low-standard deviation) and a wider (higher-standard deviation) Gaussian signals in the real space, representing a particle and a reflection, respectively. When the Fourier transform is applied, the reflection signal appears as a narrower Gaussian mainly represented by smaller wavenumbers, while the particle signal appears as a wider Gaussian with higher associated wavenumbers. Recall from Section 2.3.2.1 that the 0 wavenumber (also called

DC component) represents the mean of the signal, so when dealing with images, this component will contain the average of the whole image intensity. Therefore, as reflections normally dominate the image, the reflection content in Fourier domain is expected to be concentrated on the lowest wavenumbers, close to this DC component.

Additionally, a unit impulse signal is also presented in Figure 4.2 (e). This serves as an example of the case when particles are too small or too far away that the imaging system is not able to capture their whole shape. When the Fourier decomposition is applied, it results that this signal contains all single-wavenumber components with unit magnitude. This means that to reconstruct the unit impulse signal, all the frequencies need to be taken into account.

4.1.1. Working principle

Taking advantage of the difference in wavelengths between reflections and particles, the spatial Fourier transform (presented in Section 2.3) can be used to decompose the image signal into wavenumber components. Low wavenumbers will correspond to reflections, while high wavenumbers will correspond to small objects, like particles. The key to distinguishing between reflections and particles lies in filtering out the low wavenumbers that correspond to reflections and keeping the particle images by leaving unaffected the content corresponding to high wavenumbers. This can be done by employing a high-pass filter on the Fourier transform signal to modify the image content information. By doing so, the unwanted large-scale intensity variations from the image data are removed, allowing to isolate the signal associated with particles. In Fourier analysis for image processing, there are three main steps: obtaining the Fourier transform, modifying it and performing the inverse Fourier transform operation. Knowing this about Fourier analysis, the following pipeline is introduced as proposed methodology (check the code for this approach in Appendix A.1). A more detailed examination of each step is presented below.

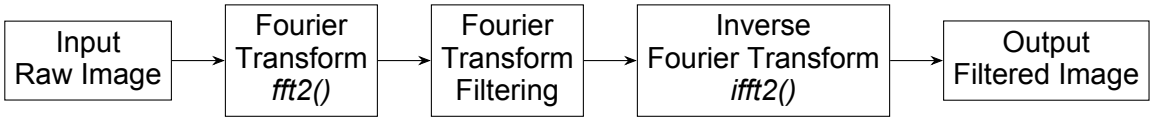

4.1.1.1. Fourier transform

Following the sphere case example, take a raw image in Figure 4.3 (a) and perform the 2D Fast Fourier transform (*2dftt()* in Matlab) operation, obtaining Figure 4.3 (b).

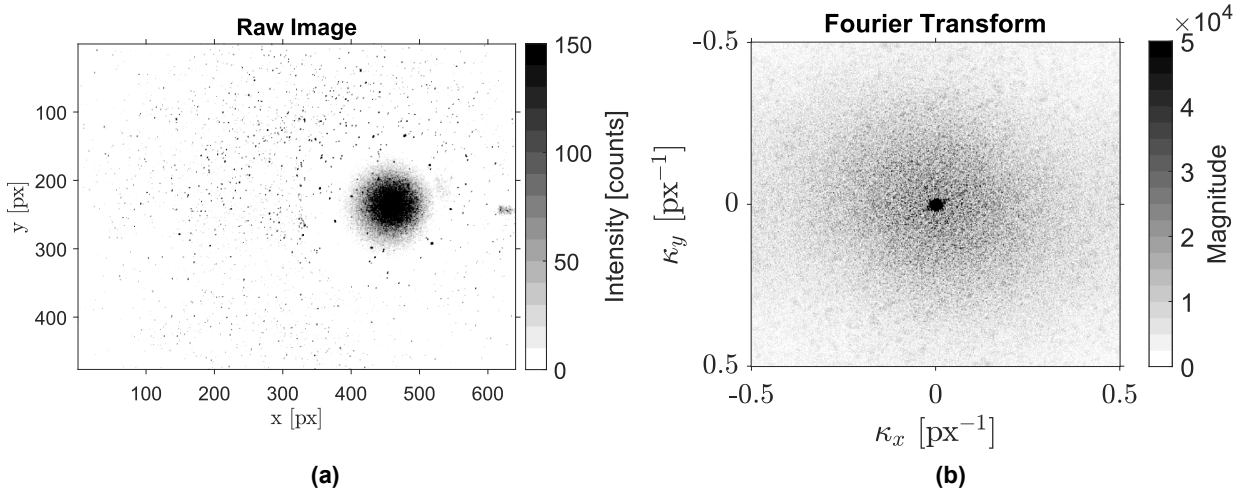

Figure 4.3: (a) Raw image and its (b) Fourier transform.

Recall that the 2D Fourier Transform of an image decomposes the image into its constituent spatial frequencies, revealing information about the variation of pixel values across the image. Low-frequency components correspond to large-scale features with smooth variations in pixel values across the image and are usually located near the center of the Fourier transform image (DC component). High-frequency components represent rapid changes in pixel values, such as edges, textures, and fine details in the image. The magnitude of the Fourier transform at a particular point indicates the strength of the corresponding spatial frequency component. Higher magnitudes indicate a stronger presence of that frequency in the image. The intensity distribution in the wavenumber space features a broad distribution with a grainy pattern. The latter is due to the random distribution of particles. A small peak at the origin (DC) is noticed in the figure, which is due to the pixel intensity being semi-definite positive, yielding a nonzero image mean value.

4.1.1.2. Fourier transform filtering

To eliminate the unwanted reflection, the high-pass filter of Equation 4.2 is applied.

$$H(\kappa_x, \kappa_y) = 1 - e^{-\frac{\kappa_x^2}{2\kappa_f^2}} \cdot e^{-\frac{\kappa_y^2}{2\kappa_f^2}} \tag{4.2}$$

This filter takes the form of a 2D Gaussian distribution centered at the 0 wavenumber component, with its standard deviation determined by the cut-on wavenumber κ_f. This shape of the filter is selected to avoid the Gibbs effect and have smooth transitions in the filtered Fourier transform, ensuring a good image reconstruction. The cut-on wavenumber plays a crucial role in this process, and for that, three criteria must be satisfied. First, the *criterion of non-rejection of particle intensity* states that the cut-on wavenumber should be smaller than the inverse of the particle image diameter to ensure that particles are not removed. Secondly, *criterion of rejection of reflection intensity* defines that κ_f should be larger than the inverse of the length of the reflection, guaranteeing the removal of the reflection. Last but not least, the *criterion of separability* ensures that the size of the reflection is significantly larger than that of the particles, creating a distinct range of wavenumbers between them. This ensures a clear separation (as depicted in the figure below) and prevents unintentional removal of particles while eliminating the reflection.

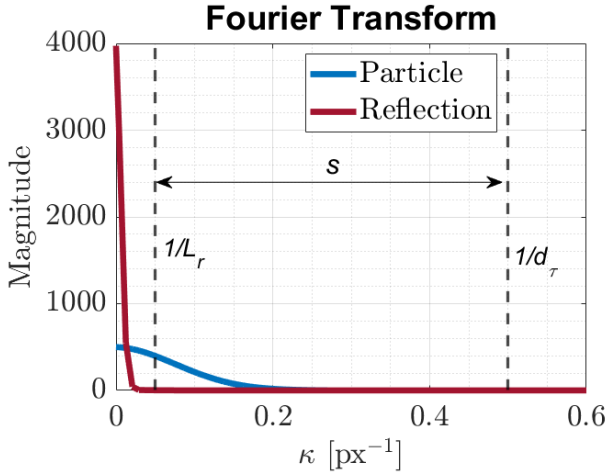

Criterion of non-rejection of particle intensity

$$\kappa_f < \frac{1}{d_\tau}$$

Criterion of rejection of reflection intensity

$$\kappa_f > \frac{1}{L_r}$$

Criterion of separability

$$s = \frac{L_r}{d_\tau} >> 1$$

Figure 4.4: Fourier Transform magnitude for particle (blue) and reflection (red) signals.

Multiplying the Fourier-transformed image with the high-pass filter allows to obtain a filtered version of the original input image by employing the inverse Fourier transform operation. A series of filter examples, along with the resulting filtered Fourier transform and filtered images, are presented in Figure 4.5. These examples demonstrate the effect of the cut-on wavenumber κ_f on the Fourier transform and, consequently, on the filtered image.

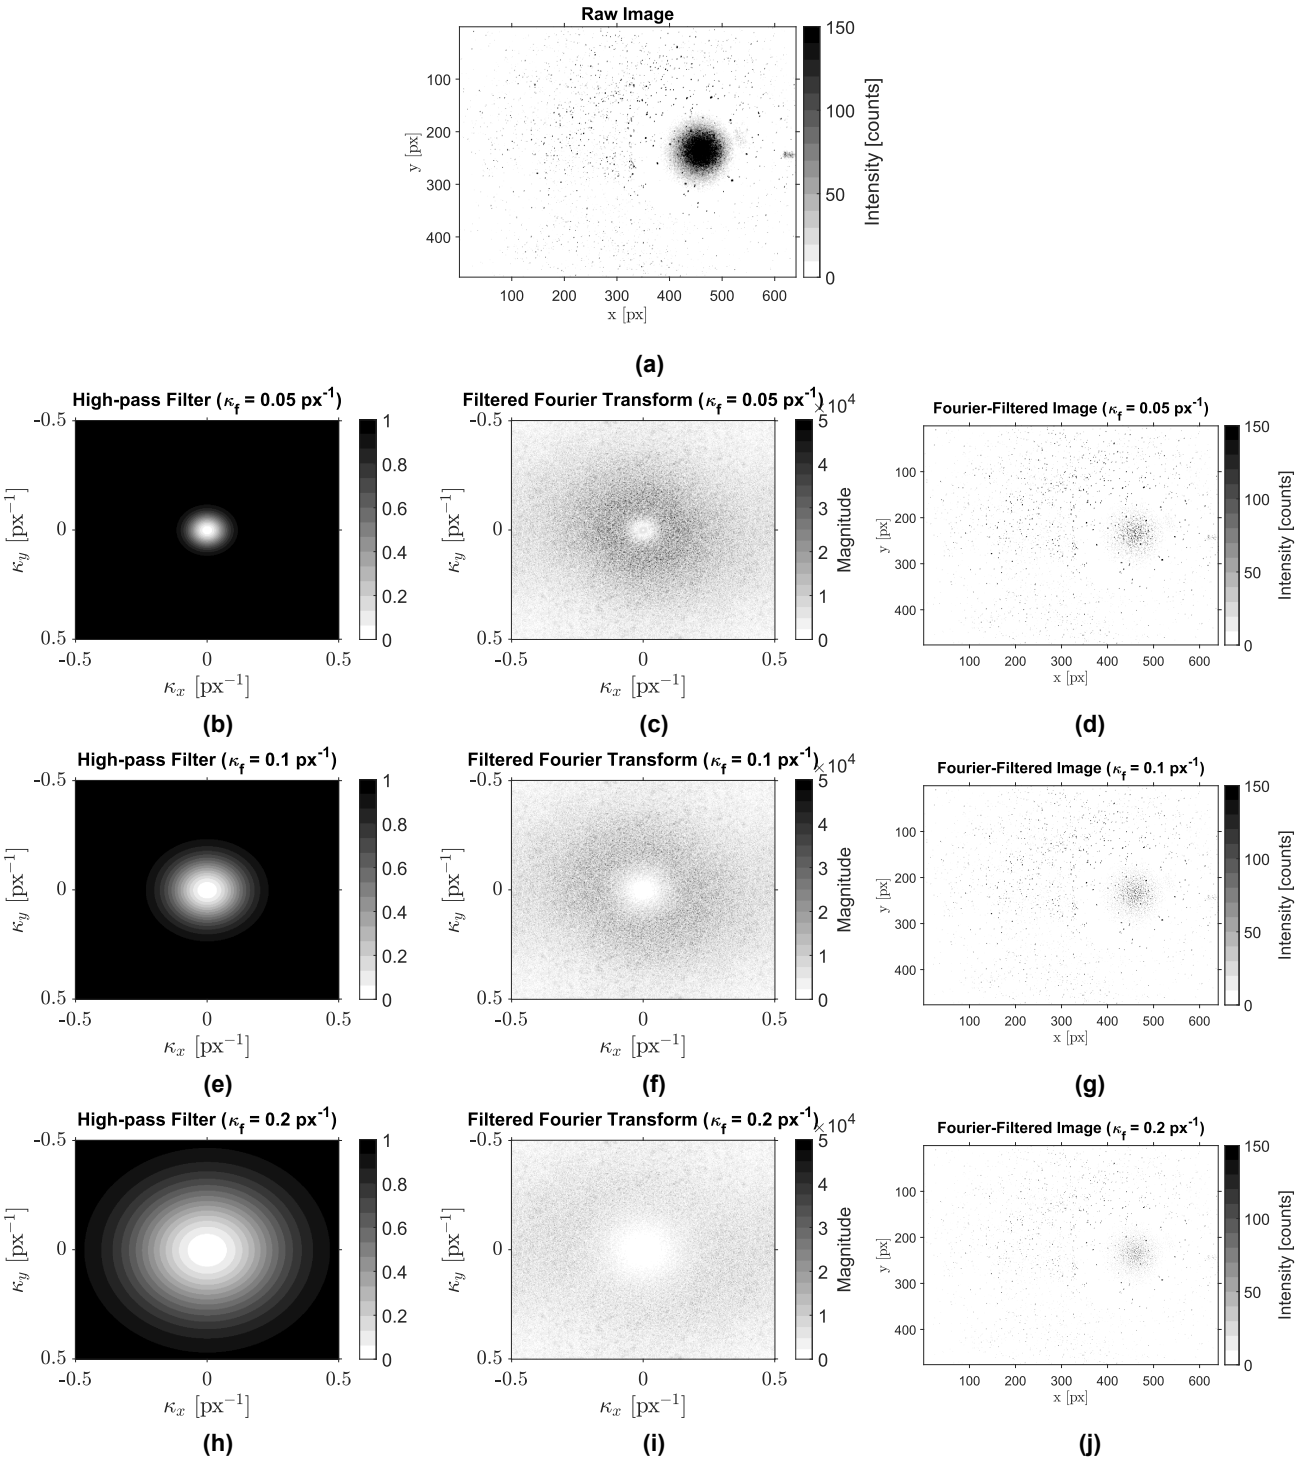

Figure 4.5: (a) Raw image. (b)-(j) Left: high-pass filters with different cut-on wavenumbers. Middle: filtered Fourier transforms. Right: resulting filtered images.

To check whether the resulting filtered images are well suited, the Signal-to-Noise ratio or SNR can be checked as image quality parameter. This SNR is typically used in PIV to check whether particle signal intensity versus the reflection or noise in the image is sufficient to track the particles. This ratio should be as high as possible, meaning that the particles' intensity is utterly dominant over the reflection. Here, the SNR is defined as the relation of the particle intensity over the reflection intensity, where the particle intensity is computed by taking the 100 largest local maximums of a region with only particles and the reflection intensity is considered

by averaging the region with only the reflection. Figure 4.6 shows a plot of the SNR vs different cut-on wavenumbers for the sphere case previously presented.

$$SNR = \frac{\text{particles intensity}}{\text{reflection intensity}} \tag{4.3}$$

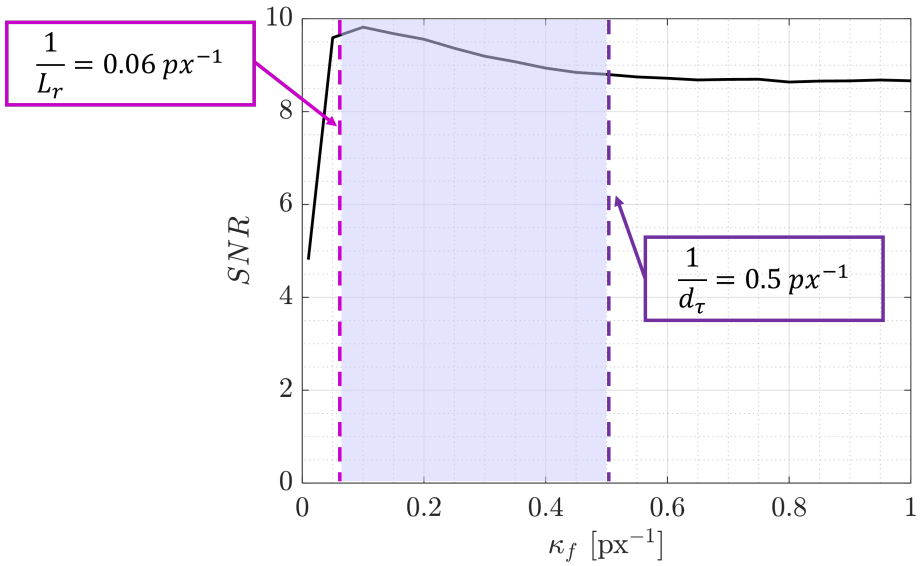

Figure 4.6: Signal-To-Noise ratio (SNR) vs. κ_f for the sphere case.

As κ_f increases, the size of the filter increases, leading to greater attenuation of image intensity in the spatial domain. The SNR correspondingly increases, suggesting that higher κ_f values may produce more desirable outcomes. However, while a larger κ_f may yield an acceptable SNR, when transformed back to the space domain, the resulting image exhibits a reduction of the reflection intensity but also some of the particles, yielding distortions to the shape of particles. This distortion can deform particles, causing them to lose their original shapes and, thus affect their proper tracking. Figure 4.7 shows this effect on a region with only particles for the three filter sizes studied previously.

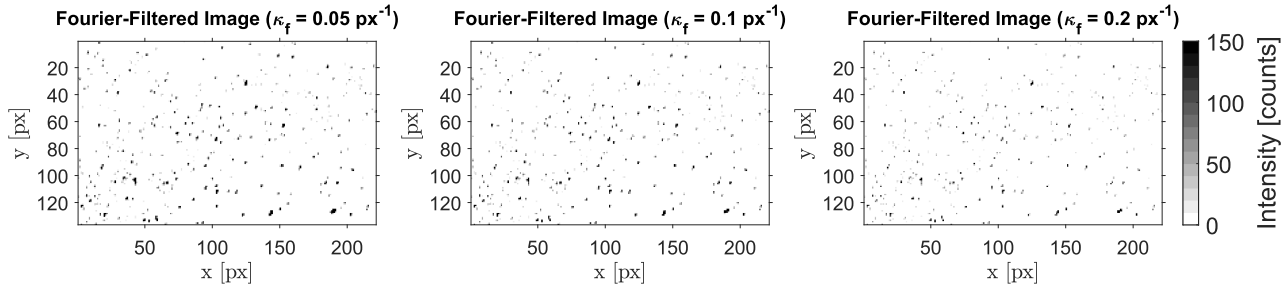

Figure 4.7: Signal-To-Noise ratio (SNR) vs. κ_f for the sphere case.

A κ_f of 0.2 px^{-1} means that the intensity of all objects with wavelengths larger than 5 px will be attenuated. Hence, even smaller particles will suffer this attenuation as the filter size increases, even leading to particles losing their original shape. This effect is more notable in larger particles as their wavenumber content is closer to the cut-on wavenumber.

To quantify this effect, the distance $\Delta\hat{\epsilon}_G$ between particle peaks of the raw and Fourier-filtered images with sub-pixel accuracy is computed. Figure 4.8 shows 150pixel x100 pixel region from both raw and Fourier- filtered images where the sub-pixel position of the particles will be evaluated.

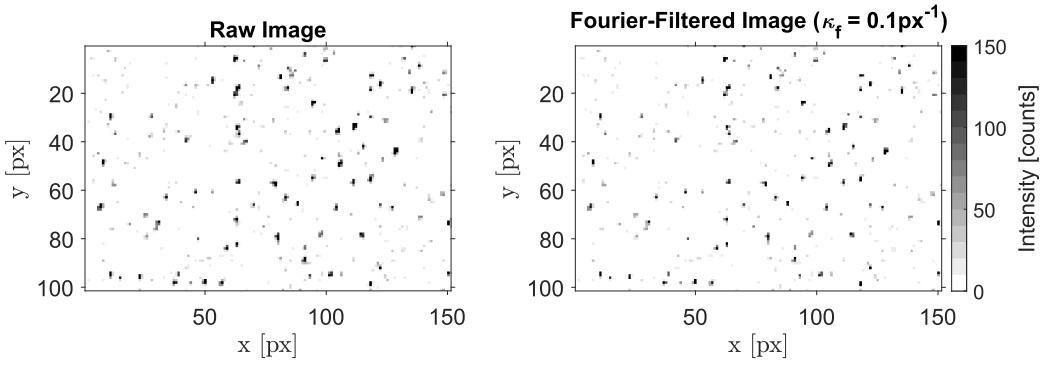

Figure 4.8: Region within (left) raw and (right) Fourier-filtered images.

This sub-pixel position $\hat{\epsilon}_G$ is obtained using Equation 4.4 [87], where R_0^* is the intensity of the central pixel, R_{-1}^* the left pixel and R_{+1}^* the right pixel (or bottom and top if considering the vertical position).

$$\hat{\epsilon}_G = \frac{\ln R_{-1}^* - \ln R_{+1}^*}{2\left(\ln R_{-1}^* + \ln R_{+1}^* - 2\ln R_0^*\right)} \tag{4.4}$$

Once the distances between raw and filtered images $\Delta\hat{\epsilon}_G$ are computed, the probability density function (pdf) of these values shown in Figure 4.9 is obtained. The mean and standard deviation of the resulting distribution is 0 px and 0.05 px, respectively. This indicates that the effect in the particle intensity and shape due to the implementation of the Fourier filter is significantly low; particle peaks suffer almost no change when applying the filter.

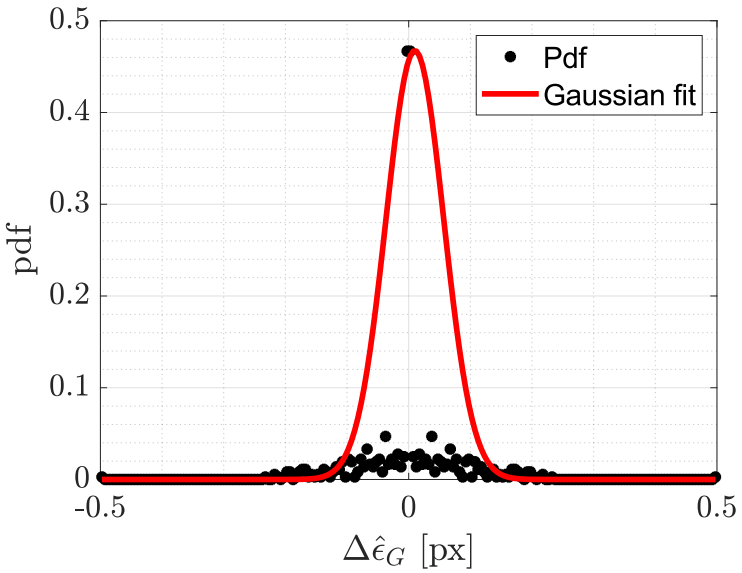

Figure 4.9: Probability density function of $\Delta\hat{\epsilon}_G$ (between raw and Fourier-filtered images).

The user is advised to make the selection of the filter's cut-on wavenumber κ_f based on the criteria presented above and the overall size of the particles on the particular case of interest.

4.1.2. Results on CVV measurements for the flow over a sphere

Let us consider the filter with κ_f = 0.1 px^{-1}. The resulting filtered image shows a significant attenuation of the reflection compared to the original image; however, residual components of the reflection remain visible.

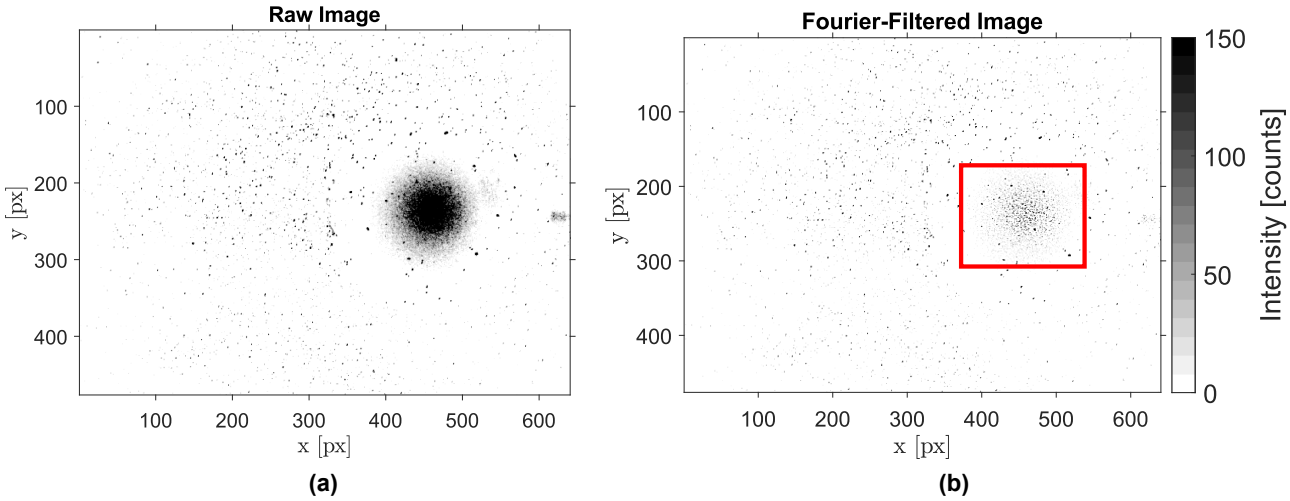

Figure 4.10: (a) Raw image and its (b) Fourier-filtered image.

Upon closer inspection, it becomes clear that within the area of the reflection (denoted in red in the previous figure) in Figure 4.11, particles are present, but also other high-wavenumber components that remain unaffected by the filter. This is because light reflections in experimental images are non-Gaussian and often contain a bit of granularity (represented by smaller wavelengths) due to surface imperfections.

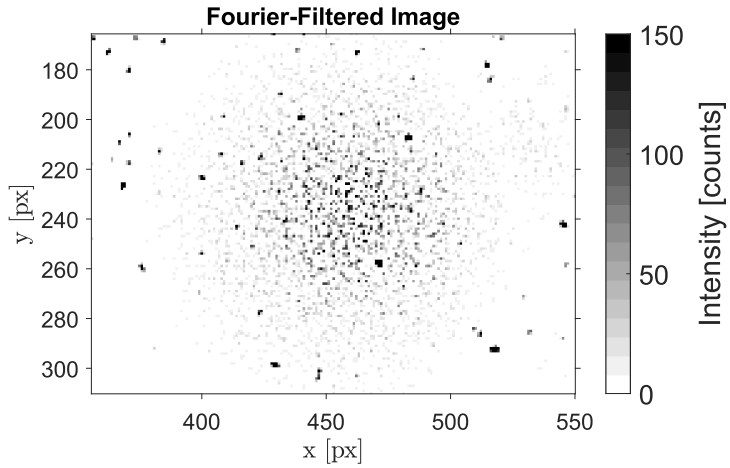

Figure 4.11: Close-up on the reflection region of the Fourier filtered image ($\kappa_f = 0.1\ px^{-1}$).

Figure 4.19 shows the resulting velocity field after processing 100 images with the Shake-the-Box algorithm of three different cases: raw images (no image pre-processing applied), Butterworth time-filter and Spatial Fourier Filter.

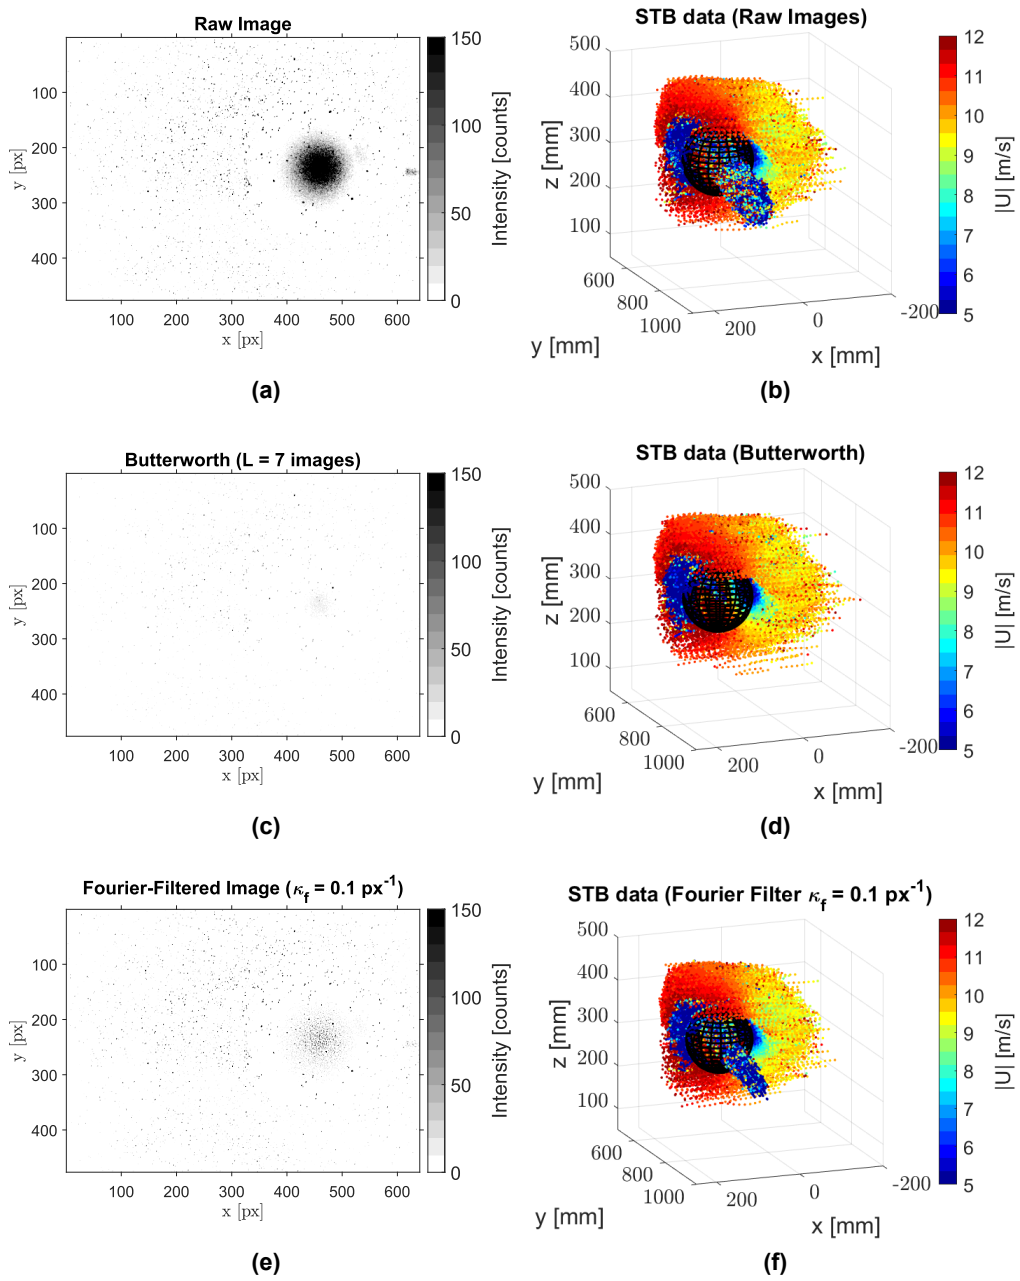

Figure 4.12: Left: Processed images (from top to bottom: raw images, time-filtered Butterworth and spatial Fourier-filtered). Right: Resulting velocity field.

No image pre-processing clearly affects the STB data as the algorithm tracks particle tracers within the reflection region that in reality do not exist. On the contrary, the Butterworth time filter effectively removes the reflection due to its steadiness over time, resulting in a clean-of-reflections velocity field. The Spatial Fourier Filter strategy reduces the reflection's presence, resulting in an improvement compared to the untreated images case. However, particles are still visible within the region affected by reflection as an accumulation of particles that propagate along the camera's line-of-sight, albeit smaller in size compared to when processing the raw images. Despite this small improvement, it can still impact the accuracy of results.

As a summary, this filtering strategy attenuates the reflection intensity compared to that of the particle tracers. However, as the reflection contains high-wavenumber components, the choice of the filter size κ_f may result in the reflections not being fully removed or the particle images being attenuated and distorted by the filter. Besides, light reflections in experimental images do not follow exactly a Gaussian distribution and often contain certain level of granularity or speckle (represented by smaller wavelengths) due to surface imperfections and laser coherence. Therefore, to overcome this limitation, the methodology presented in the following chapter is introduced.

4.2. Spatial frequency-based approach for reflection identification and masking

In order to overcome the limitations posed by the previous reflection treatment approach, a second methodology is introduced in this chapter. The objective of this method is to completely eliminate the regions of reflections. To do this, two main steps will be followed: first, identify the reflection(s) and the corresponding location and secondly, create a mask and apply it to the original image to remove the reflection.

4.2.1. Working principle

In order to mask the reflection specifically, identify and locate it first is a must. It is clear that is easy to do this by human eye, however when large datasets of images need to be processed, there must be a way for the computer to automatically do it for every image (see Appendix A.2 for the implementation of the code). Recall the filtered image obtained from the Fourier-filtering method introduced in the previous chapter in Figure 4.13 below.

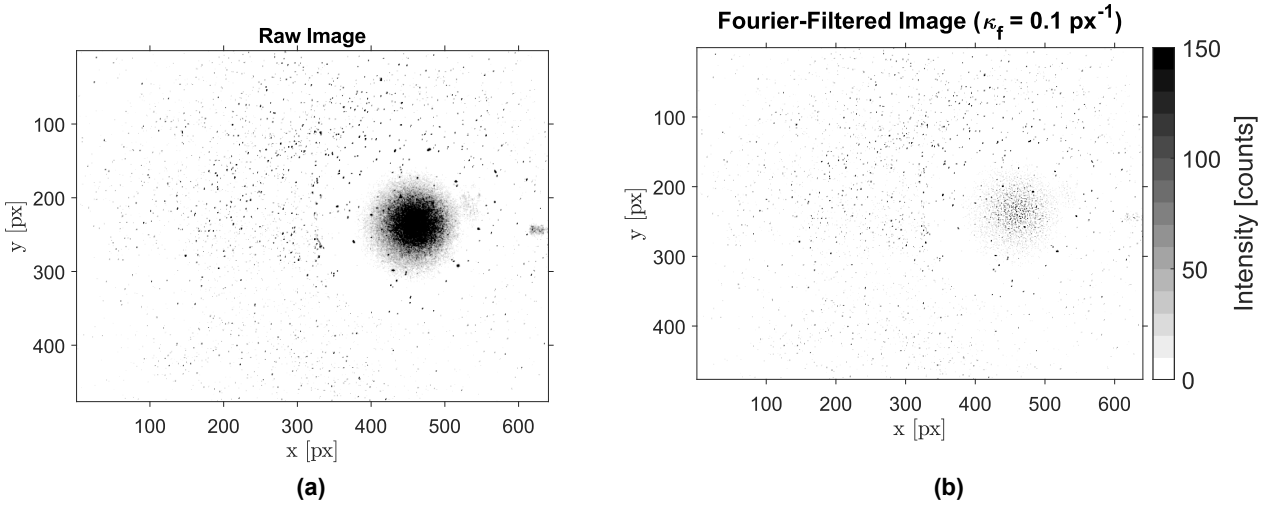

Figure 4.13: (a) Raw image and its (b) spatial Fourier-filtered image.

There is a clear difference between the raw image and the filtered one, particularly in the region dominated by the reflection. The reason is that the high-pass filter is known to affect the intensity of the reflection more than that of the tracer particles. Therefore, an option to highlight and locate the reflection(s) is to quantify which parts of the raw image have changed most when the spatial Fourier high-pass filter has been applied. This is made using the Structure Similarity Index Measure (SSIM), which is a widely used method for measuring the similarity between two images [88]. SSIM compares local patterns of pixel intensities between the two images being

compared. The similarity is evaluated in terms of luminance, contrast, and structure, which are key components of human perception – see Figure 4.14. Below there is a brief summary of these three components:

- **Luminance Comparison**. SSIM measures the similarity in terms of brightness between corresponding pixels in the two images.
- **Contrast Comparison**. It examines the contrast similarity, which refers to the difference in brightness between neighboring pixels.
- **Structure Comparison**. SSIM evaluates the similarity in terms of image structure, which captures spatial dependencies among pixels.

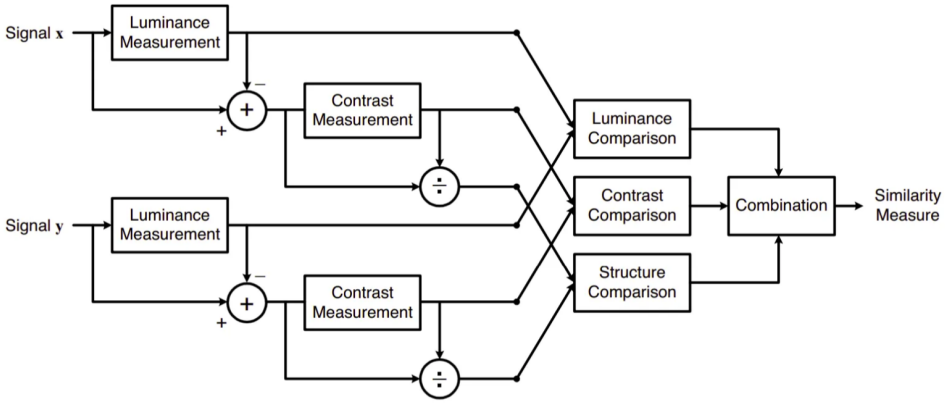

Figure 4.14: SSIM pipeline. Source: [88].

This index provides a normalized comparison between images with a range of SSIM values between -1 and 1. A value of 1 indicates perfect similarity, thus that the compared images are identical in terms of structure, contrast, and luminance. In other words, there is perfect similarity between the images. When SSIM ≈ 0 indicates little to no similarity between the images, which could mean that the images are significantly different. Last, but not least, a value of -1 indicates perfect dissimilarity. Negative values of SSIM imply that the compared images are not only dissimilar but are also inversely related in terms of structure, contrast, or luminance. In the case of the spatial Fourier filtered image, only the intensity of the image changes (while structure and contrast remain unchanged). This yields a SSIM value that may decrease but generally remains relatively high compared to more significant changes. This is because SSIM is designed to be robust to changes in intensity while primarily focusing on the structural and contrast similarities between images. Therefore, in the cases that concern this project, extremely low values of SSIM will not be expected, but rather slightly lower values in the regions that have suffered most intensity change due to the spatial Fourier filter.

Let's continue with the example of the sphere case. Figure 4.15 shows the SSIM map of the spatial Fourier-filtered image with respect the raw image (Figures 4.13 (b) and (a), respectively). A higher SSIM value suggests greater similarity between the images being compared and the low values correspond to areas affected by reflections. In the low values regions, there are also a few particles that are being picked up. These particles have suffered a decrease in their intensity after the spatial Fourier filtering, therefore this is why they appear as low SSIM value regions in the SSIM map. However, the reflection region clearly shows two main characteristics: low SSIM value and large area. In order to keep only the reflection region, two additional steps will be required.

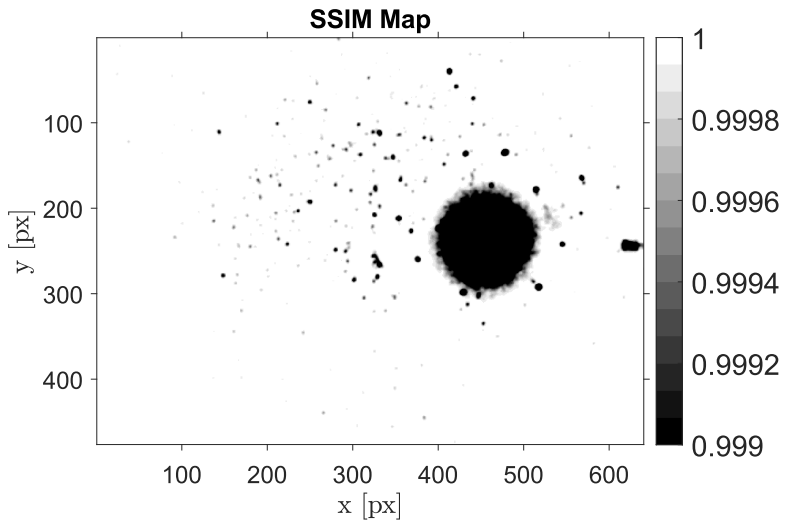

Figure 4.15: SSIM map on the sphere case image from Figure 4.13.

First, a threshold on the SSIM value is applied to the SSIM map to keep only the regions that have changed most (assuming that the reflection is the part of the image that would suffer the largest variation). For instance, in SSIM map presented above, the reflection region certainly contain values of SSIM below or equal to 0.999 (called $0.(9)_3$ henceforth). The mask that results from applying this values as cut-on threshold (everything below or equal this value will be 0 and above 1) is presented in Figure 4.16.

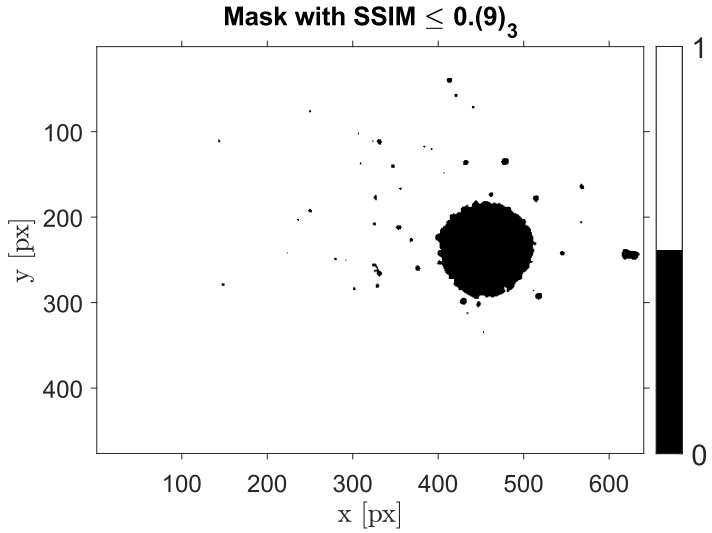

Figure 4.16: Mask obtained after applying SSIM threshold to SSIM Map.

As shown in Section 3.1, particles are composed of high wavenumber components, but also low wavenumber content. When applying the high pass filter, the intensity of these particles inevitably decreases, leading some of them to exhibit low SSIM values in the SSIM map, similar to reflections. Including particles in the mask can result in data loss, as the algorithm cannot accurately track them. To specifically retain shapes associated with reflections, the area of each blob in the mask is computed. Particles typically have a diameter of 3 to 10 pixels and assuming

that after applying the SSIM threshold their area can slightly increase, an area threshold of 200 pixels is used in this case. Then, only the shapes with an area greater than this value are retained as shown in Figure 4.17.

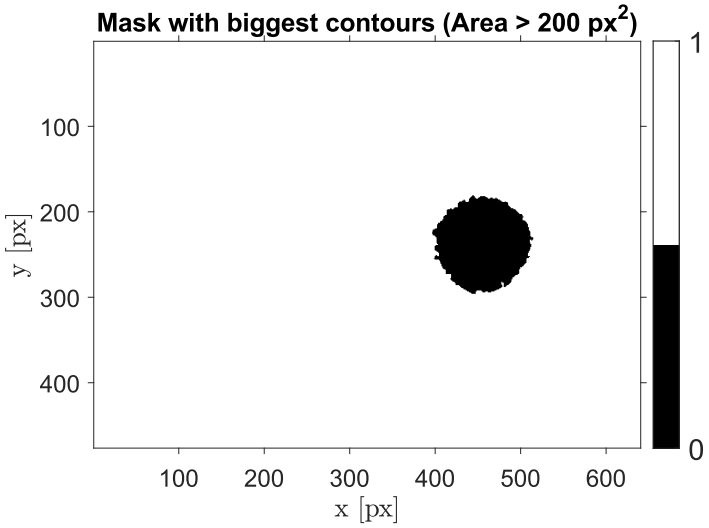

Figure 4.17: Mask obtained from considering the blobs with area larger than 200 px^2.

The resulting mask is then applied to the spatial Fourier-filtered image to obtain the masked image in Figure 4.18 without the reflection.

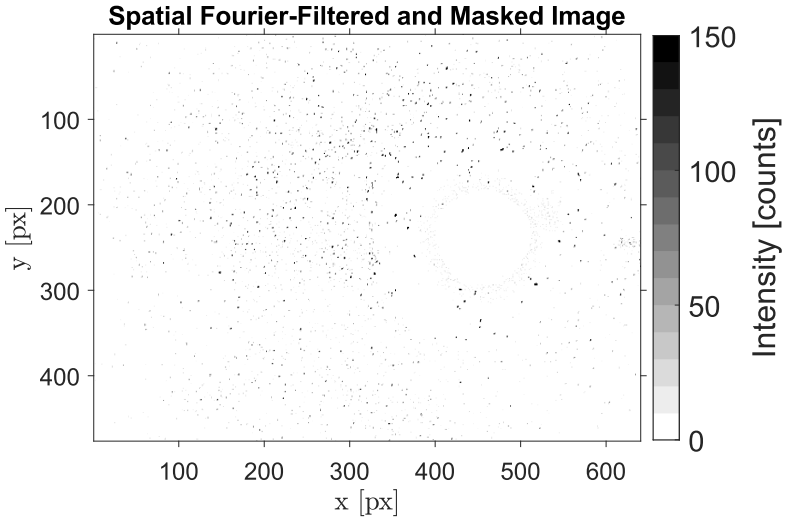

Figure 4.18: Masked filtered image.

Now let's compare this methodology with other image-processing techniques: no filter applied (raw images), Butterworth time filter and Spatial Fourier Filter. Figure 4.19 shows the pre-processed images on the left and the resulting Shake-the-Box data on the right.

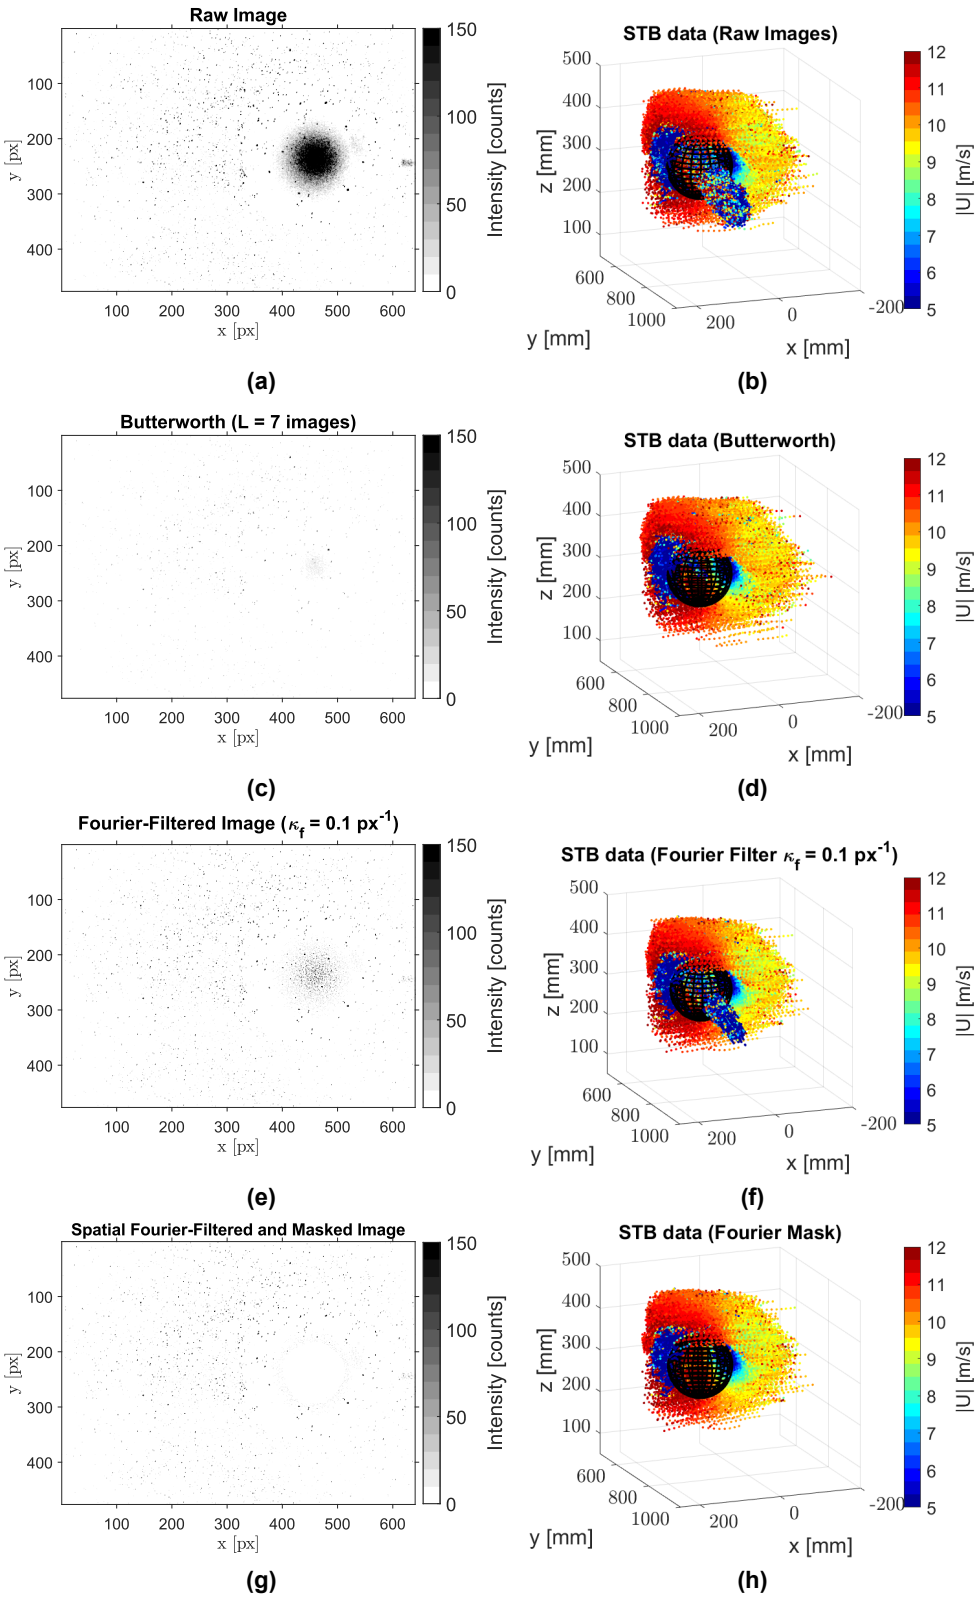

Figure 4.19: Left: raw, Butterworth time-filtered, spatial Fourier filtered and Fourier masked images (top to bottom). Right: resulting Shake-the-Box data.

The comparison of the four cases shows distinct outcomes in treating light reflections. As discussed previously, both untreated and Spatial Fourier Filter cases show particles within the reflection region, although the latter approach reduces the size of this area. On the other side,

the Butterworth time filter accurately removes the reflection, resulting in a clean-of-reflections velocity field (Figure 4.19 (c-d)).

The Spatial Fourier Filter + Mask approach shows a void of spurious regions due to the reflection (Figure 4.19 (g-h)), and confirms that the image pre-processing technique can robustly eliminate the contribution of light reflection areas while leaving the rest of the domain unmodified. Due to the masking operation, possible particles going over these regions will also be masked. Thus, regions in space along the lines of sight of affected cameras will be empty of particle trajectories. However, for the case of robotic volumetric PIV, these can be measured and filled in from a different robot position. Therefore, the previously presented Spatial Fourier Filter procedure used to filter out the reflection can be made less sensitive to the choice of κ_f if the filter is not used directly as a weighting function, but rather to guide the operation of automated masking.

4.3. 3D Particle Concentration-based reflection masking

The reflection treatment techniques proposed in the previous sections involve its implementation in the image pre-processing stage. However, in the case that image pre-processing techniques are not successful in properly removing reflections, these would still appear as erroneous data in the Shake-the-Box results corresponding to not-fully removed reflections. Therefore, a reflection treatment method should be applied in a later stage of the Robotic PIV processing pipeline. This section introduces a third reflection treatment approach that aims to remove any residual region of reflections in the Shake-the-Box data by particle concentration analysis.

4.3.1. Working principle

Having characterized reflections in STB data as accumulations of particles that propagate along the line-of-sight, a methodology that analyses the particle concentration on each bin of the measurement volumes is proposed. This relies on the characteristic described in Section 3.2 that reflections will generally appear in STB data as regions of high particle concentration compared to the real tracks. Hence, the aim of this method is to isolate regions of high particle concentration that theoretically belong to spurious regions (e.g. reflections), and with them create a mask to apply to the instantaneous STB data. With this principle in mind, the following steps are proposed.

1 Shake-the-Box track data

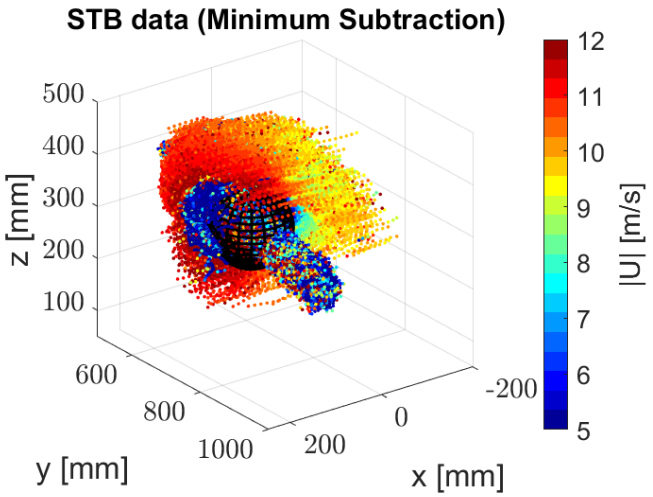

Figure 4.20: Shake-the-Box data on the pre-processed images with Minimum Subtraction (over entire series) for the sphere case.

The process begins with the Shake-the-Box results. Following the trend of the previous chapters, the explanation of this method will also use the sphere case data. Recall the Shake-the-Box results from processing the pre-processed images with Minimum Subtraction (over entire series) in Figure 4.20. In the 3D results, there is a clear spurious region that crosses the sphere that does not convey any physical meaning of the flow field, and should therefore be removed to prevent misinterpretation.

2 Binning and average particle concentration distribution

Once obtained the STB data, a binning step is performed over the measurement volume. In this particular case of the sphere, only one measurement volume is analyzed. However, in scenarios where multiple measurement volumes are present, the current method should be applied to each of them individually. Recall that the binning divides the volume into small cells or bins, and averages the flow field information over time. Figure 4.21 shows the particle concentration distribution over the domain around the sphere after performing this step. As stated in Section 3.2, the region due to the reflection shows a higher values of particle concentration compared to the rest of the flow domain.

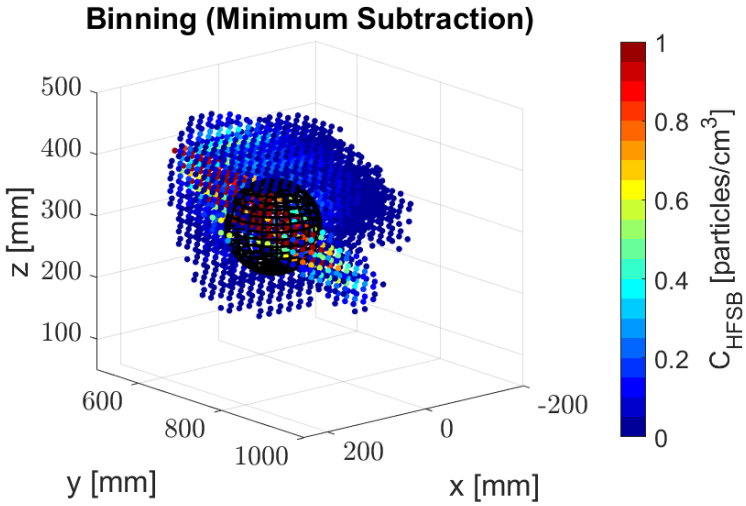

Figure 4.21: Particle concentration C_{HFSB} on the pre-processed images with Minimum Subtraction (over entire series) for the sphere case.

Therefore, there is a way to separate good regions of real particles (exhibit low C_{HFSB} values) from wrong regions of reflections (higher C_{HFSB}) by applying a threshold on the particle concentration distribution.

3 Particle concentration reference value

The particle concentration threshold or reference value C_{HFSB_0} should be such that takes into account the experimental conditions since every experiment is different. This value can be established by analysing the expected or theoretical particle concentration on the measurement region (or test section), which can be approximated as [89]:

$$C_{HFSB} = \frac{\dot{N}}{\frac{A_{rake}}{n} \cdot U_\infty} \tag{4.5}$$

where

- $\dot{N}$: effective bubble production rate of the seeding rake $\left[\frac{particle}{s}\right]$

- A_{rake}: area occupied by the seeding rake $[cm^2]$
- n: wind tunnel contraction ratio $[-]$

$$n = \frac{A_{settling\ chamber}}{A_{test\ section}} = \frac{A_{seeding\ rake}}{A_{seeded\ region}} \tag{4.6}$$

where $A_{settling\ chamber}$, $A_{test\ section}$, $A_{seeding\ rake}$, $A_{seeded\ region}$ are the areas shown in Figure 4.22.

- U_∞: free-stream velocity $[m/s]$

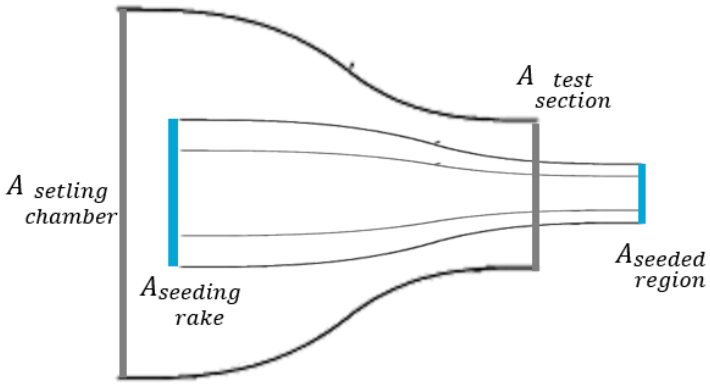

Figure 4.22: Representation of wind tunnel contraction.

For instance, considering a seeding of 30,000 bubbles/s per nozzle and assuming all nozzles are working at their 100%, the theoretical seeding particle concentration C_{HFSB} is plotted vs. the wind velocity in Figure 4.23 for the three possible test-sections of the W-tunnel (204 nozzles) and the Open-Jet Facility (OJF) wind tunnel (398 nozzles) [90].

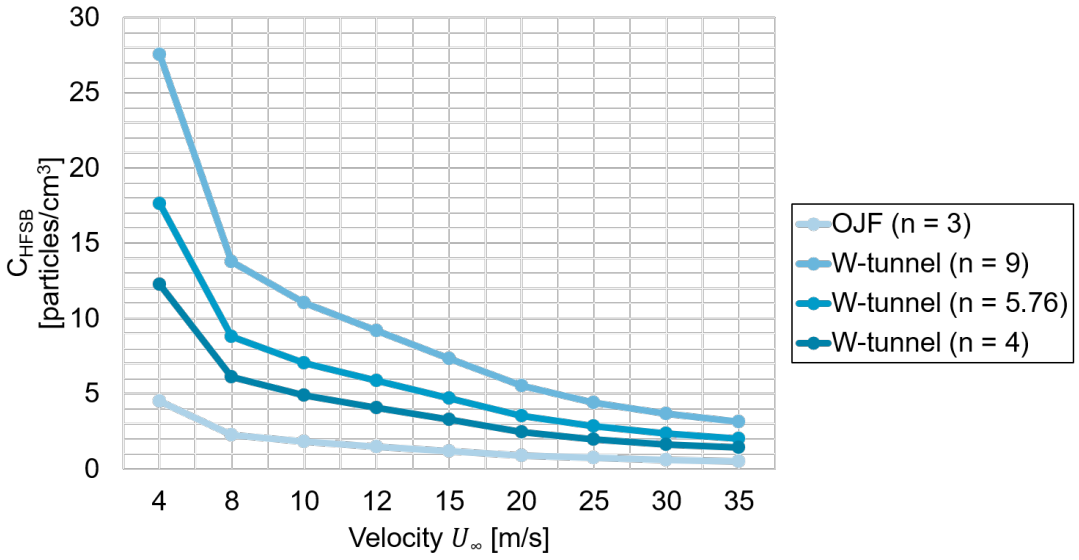

Figure 4.23: Particle seeding concentration vs. velocity for different contraction ratios and wind tunnels (considering a production rate of 30,000 bubbles/s per nozzle).

These values can be taken into consideration as reference, however occasionally they can be slightly optimistic; the seeding rake does not always work at its 100%. However,

this value can taken such that if there is any bin with a particle concentration above that value, that will certainly be a spurious region. For this reason, the user is advised to take a clean region (for instance, in the free-stream) without any artifacts, check the particle concentration in this area and consider that value as reference.

4 Mask generation from reflection data

When the particle concentration reference value C_{HFSB_0} is selected, it is used to threshold the bin data to separate between spurious regions (bins where $C_{HFSB} > C_{HFSB_0}$) and real particle tracks (bins where $C_{HFSB} \leq C_{HFSB_0}$). This threshold is applied in the Shake-the-Box data from processing the images with the minimum subtraction filter. The isolated reflection data of the sphere case is shown in Figure 4.24.

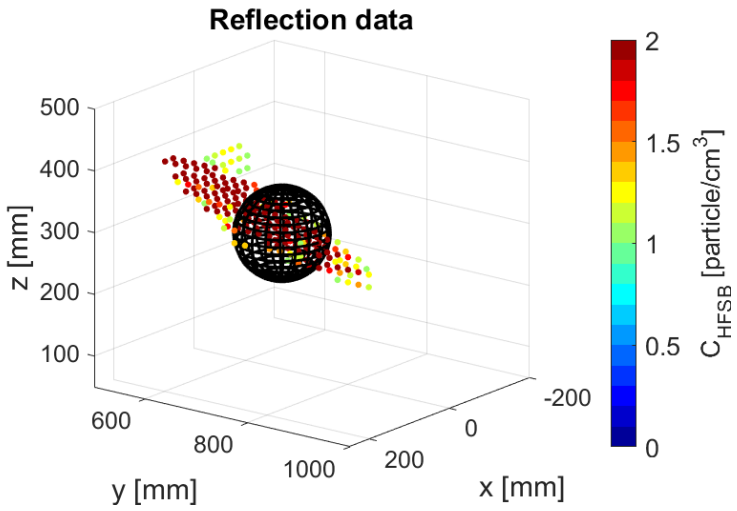

Figure 4.24: Isolated reflection binning data.

The isolated data corresponding to reflections is used to generate a 3D volume employing alpha shape approximation [91]. This is depicted in Figure 4.25. This volume is then used as a 3D mask, allowing the removal of unwanted points in the Shake-the-Box data.

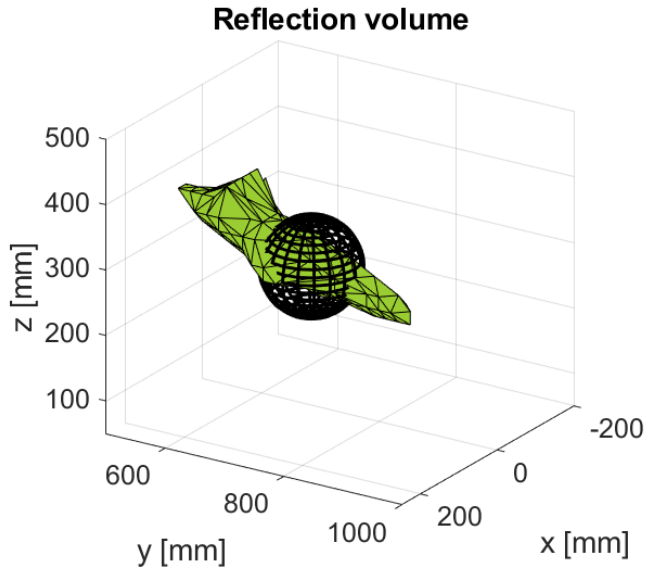

Figure 4.25: Reflection volume obtained by alpha shape approximation.

5 Resulting masked data

All points that fall inside this volume are considered reflections or artifacts and thus are removed, whereas the points outside this region are considered good tracks and are kept untouched. Figure 4.26 shows the effect of applying this mask on the Shake-the-Box data.

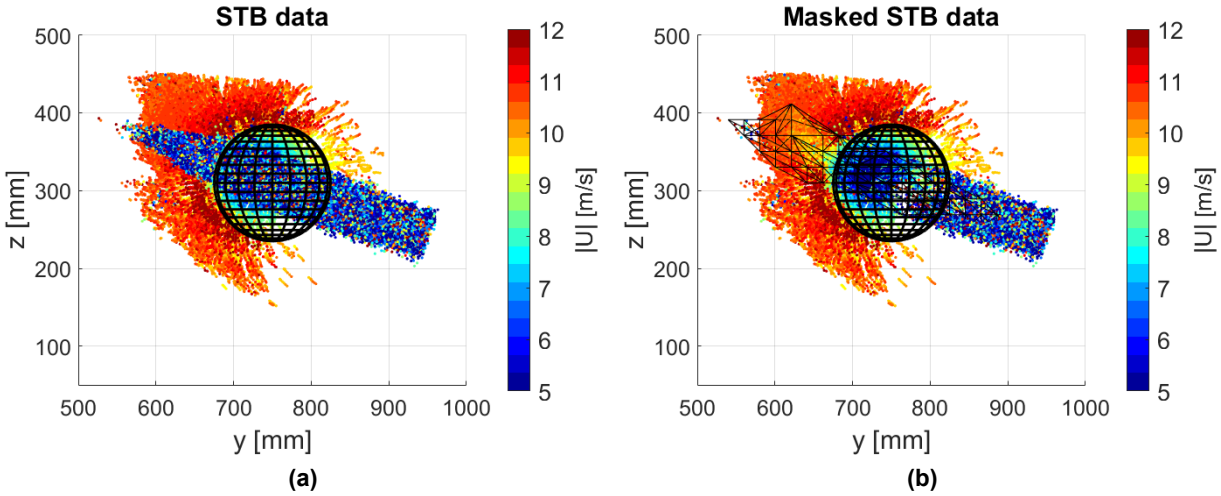

Figure 4.26: (a) Original Shake-the-Box data and (b) Masked Shake-the-Box data.

The approach is able to remove the part of the reflection that showed the largest particle concentration in the binning data. However, it is not effective at removing the rest of it due to a lower particle density in this region. Despite these preliminary results, the method will be tested on the results from the experimental campaign planned for this project.

5

Experimental setup and procedures

This chapter describes the experimental campaign performed to test out the proposed reflection treatment methodologies presented in previous chapters. The main goal of these experiments is to acquire images with regions of reflections to later apply the developed reflection treatment approaches. Therefore, different geometries that produce different types of reflections are to be tested.

5.1. Wind tunnel

The experiments were performed at the W-Tunnel at TU Delft Aerospace Engineering faculty's High-Speed Laboratory (HSL). This consists of an open jet open-return-circuit wind tunnel that allows interchangeable square exit test sections: 40×40cm, 50×50cm and 60×60cm. In this case, a square test section with a 60×60 cm^2 cross section is used (contraction ratio of 4:1). The wind tunnel allows a maximum velocity of 35 m/s, with a minimum turbulence intensity of 0.5%. The velocity is controlled by setting the revolutions per minute (rpm) of the wind tunnel fan.

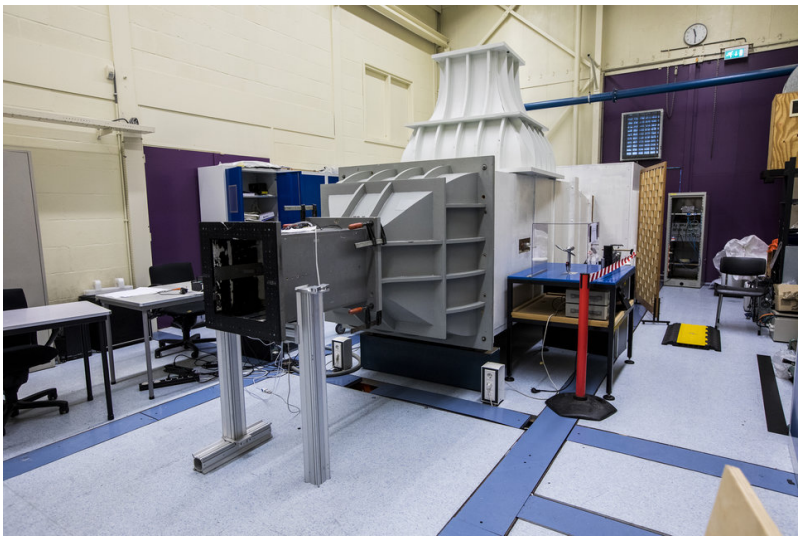

Figure 5.1: W-Tunnel at TU Delft's HSL. Source: [92].

5.2. Test models

As already stated, the main goal of these experiments is to acquire images with regions of reflections to later apply the proposed reflection treatment methodologies. For this reason, three models with different complexities in their geometry were tested during the experimental campaign. The variation in geometries will yield diverse shapes and types of reflections that will be valuable in the implementation and evaluation of the methodologies developed. Table 5.1 shows an overview of the test objects used in the experiments, which are described in more detail in the following sections.

Steady reflections		**Unsteady reflections**
Side-view mirror	Formula 1 car	Propeller
Simple geometry	Complex geometry	Moving-object

Table 5.1: Overview of models tested.

5.2.1. Side-view mirror

The first test object is a side-view mirror model, which consists of a half cylinder of 10 cm diameter with a quarter sphere attached at the top, resulting in a total height of 15 cm – see Figure 5.2 (a). This object has already been tested by Saredi et al. [93]. The model is installed on a 2.0 m long flat plate with smooth curved leading edge, which is placed at 20 cm height above the bottom edge of the exit as shown in Figure 5.2 (b).

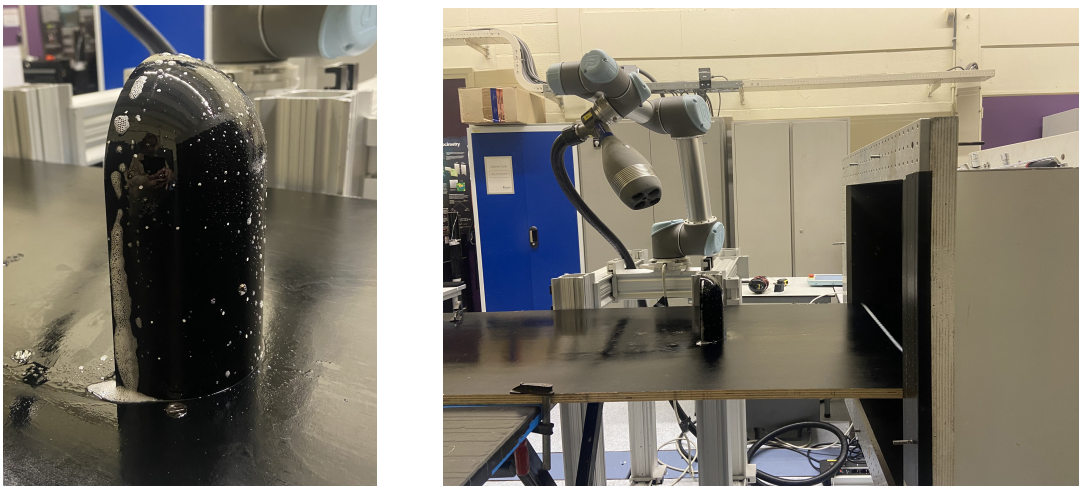

(a) **(b)**

Figure 5.2: Side-view mirror (a) model and (b) setup.

5.2.2. Formula 1 Car

The second test object is a Formula 1 car small scale model, more specifically the Fernando Alonso 2005 championship-winning Renault R25. The model is 3D-printed in-house with resin and had dimensions of $22 \times 9 \times 6$cm (length$\times$width$\times$height). This serves as a more complex geometry with all the little details that contain this type of cars, including the small side-view mirrors, wheels, front and rear wings, etc. The model is placed on the same 2.0 m long flat plate as the side-view mirror. A close-up view of the car model is presented in Figure 5.3.

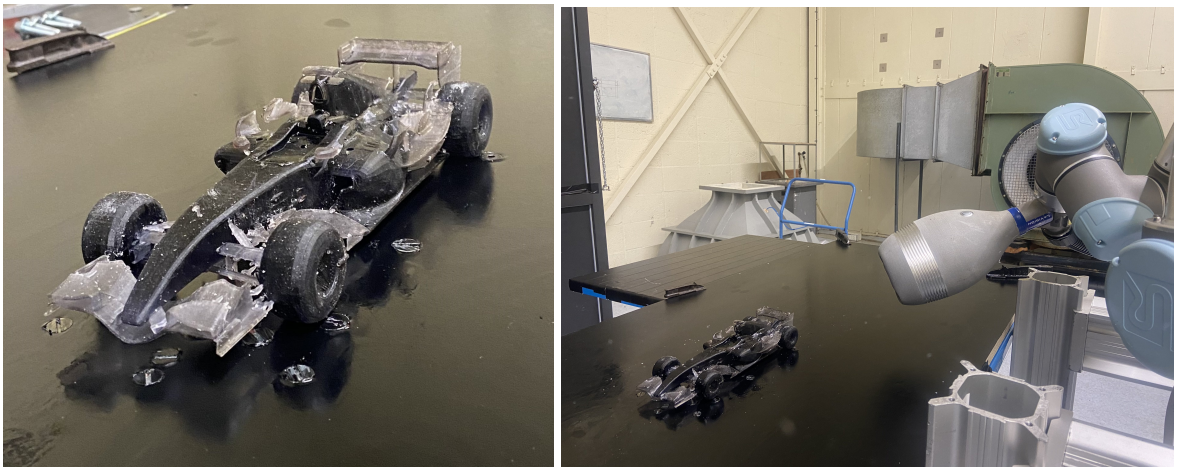

(a) (b)

Figure 5.3: Formula 1 car (a) model and (b) setup.

5.2.3. Propeller

The last object to be tested is the propeller model shown in Figure 5.4 (used and tested by Cueto in [94]). It is an APC propeller 7x5 with 2 blades and constant pitch with a diameter of 7 inch/17.7 cm installed in a pusher configuration. The propeller is powered by a Maxon Motor RE310007 60 W brushed DC motor that allows a maximum rotational speed of 9100 RPM and a maximum nominal torque of 83.5 mNm. The latter is controlled thanks to a DC power source.

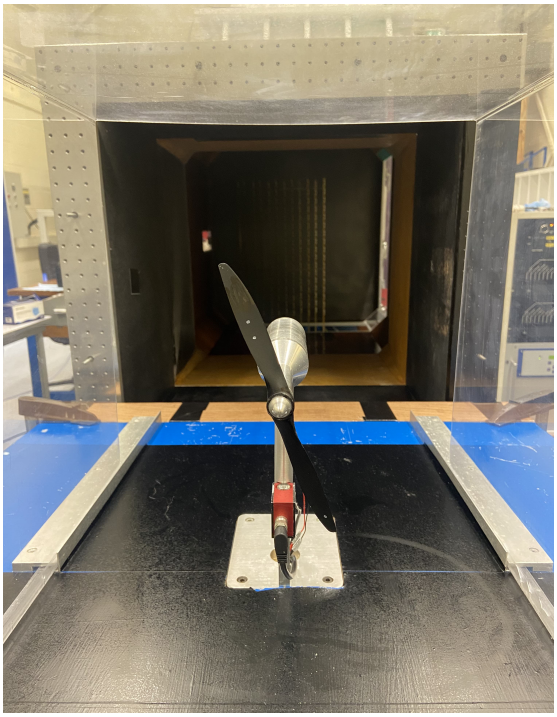

Figure 5.4: Propeller model setup.

5.3. Robotic PIV system

5.3.1. Acquisition PC

A specific PC must be used for both acquisition and processing of Robotic PIV data. This PC has the DaVis software installed. DaVis is LaVision's integrated software for data acquisition, visualization and processing. Before starting the experiments, the PC had to be properly installed and connected to the Robot and CVV hardware. DaVis version 10.2.0.74211 is used for image acquisition and processing during the experimental campaign.

To be able to control and position the robot, DaVis has to be connected to the RoboDK software. RoboDK allows the user to position the robotic arm in a virtual environment. Robot positions that have been predefined in this virtual environment can be uploaded to the robotic arm before image acquisition. The robot-RoboDK connection is achieved thanks to a router that has to be connected to the acquisition PC. Then with the robot IP address, RoboDK can detect and connect to the robot. This allows to control and move the robotic arm and set the target positions for each measurement volume to acquire. Figure 5.5 shows an example of the RoboDK interface.

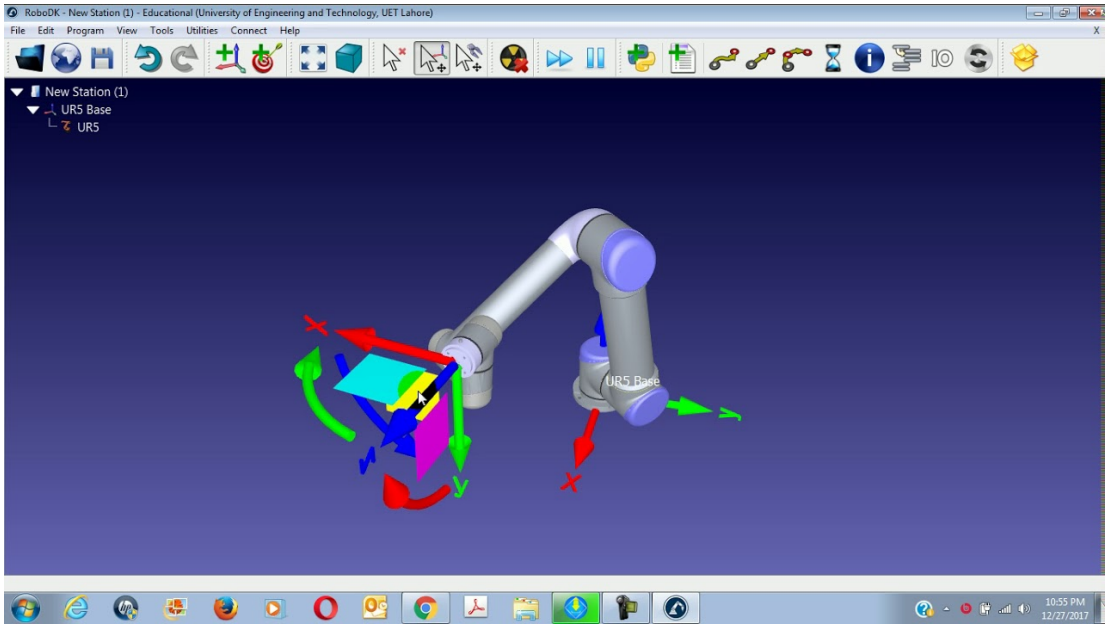

Figure 5.5: Example RoboDK interface. Source: [95].

5.3.2. Coaxial Volumetric Velocimetry system

Robotic PIV uses a coaxial volumetric velocimeter (CVV) device (see Section 5.3). Specifically, in this experiment the LaVision MiniShaker Aero is used. The CVV probe consists of four CMOS cameras (10 bits, 640×476 pixels, 4.8μm pixel pitch) and an optical fibre located between the imaging system that is responsible of the volumetric laser illumination. The laser is generated by a Quantronix Darwin Duo Nd:YLF unit (527nm wavelength, 2×25mJ pulse energy @ 1kHz), allowing the illumination in a conical shape thanks to its expansion through a spherical lens. The device is housed in an aerodynamically shaped case and is attached to the robotic arm described in the following section.

5.3.3. Robotic Arm

The CVV probe is mounted at the end of a robotic arm from Universal Robots. This is the UR5 robot, which has 6 rotating joints as degrees of freedom and 850 mm of reach [96]. Position and orientation of the arm are set either directly by the dedicated controller (robot screen) or through the interface with DaVis and RoboDK softwares. The robotic arm is installed on a X95 beam structure built on-site as shown in Figure 5.6.

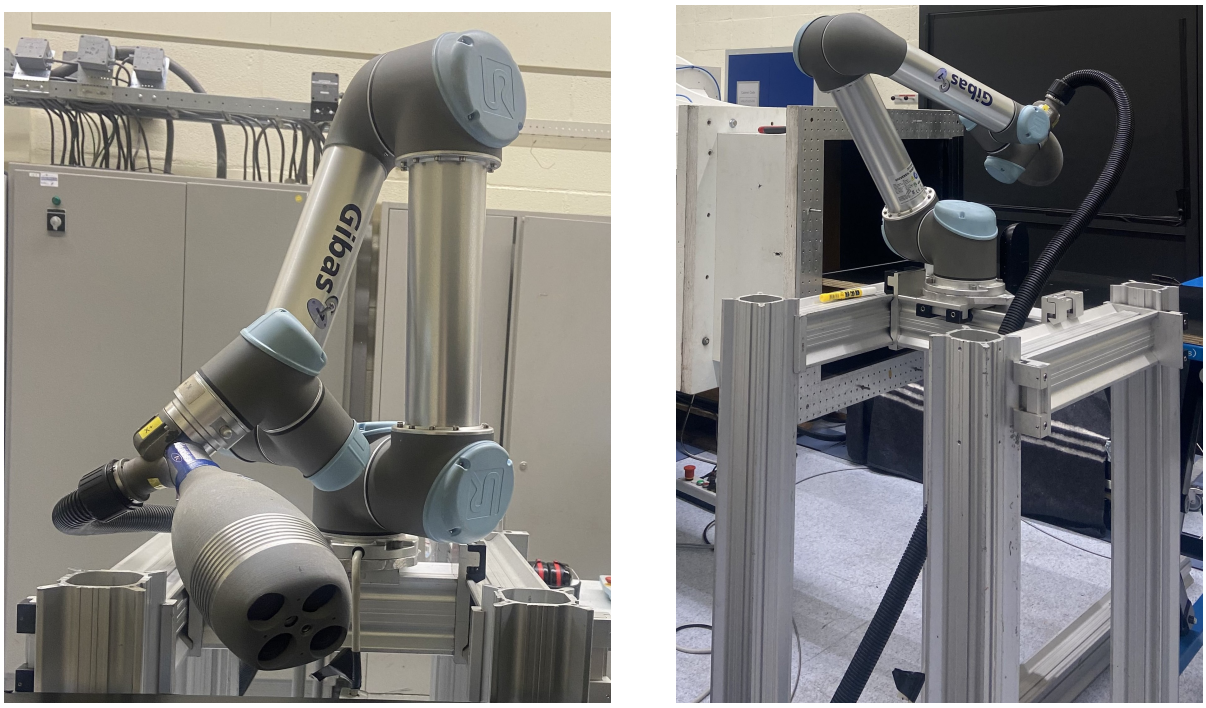

(a) (b)

Figure 5.6: (a) Robotic arm and (b) its installation setup.

The installation should not be modified or moved during the whole duration of the experimental campaign to avoid issues and affecting the system calibration. Moreover, special care has to be taken with the optical fiber mounted on the CVV probe as it is extremely fragile.

5.3.4. Seeding system

Helium-Filled Soap Bubbles (HFSB) are used with the Robotic PIV technique. To generate these bubbles, the TU Delft's HSL has an in-house built rake that consists of 204 nozzles distributed over a 10-row streamlined strut or wing assembly. Each nozzle is able to produce between 20,000-50,000 bubbles per second with a diameter of 300-500 μm [97]. This rake is placed inside the settling chamber of the W-tunnel (see behind the propeller in Figure 5.4). The seeding system includes a Fluid Supply Unit (FSU) that allows to independently modify the values of the pressure of the air, soap and helium. The effective production rate of the whole system depends on the functioning of each nozzle; it is important to take into consideration that nozzles are prone to blockage from soap accumulation or dirt.

5.3.5. System calibration

Geometric Calibration

Volumetric measurements based on multi-camera systems require a proper geometric calibration. This is the process that accurately determines the relationship between image coordinates (pixels) and real-world spatial coordinates (typically in millimeters or meters). This calibration is essential for translating the displacement of particles in the captured images into actual flow velocities. One of the most used calibration models is the pinhole calibration, which is a computer model that allows to map real-world points to the camera sensor. In this process, the parameters of the imaging system are determined and adjusted to ensure accurate and consistent measurements of objects in the real world. This calibration involves various parameters such as focal length, lens distortion, and image sensor characteristics.

The calibration process involves scanning a calibration plate in different positions: frontal view and then moving the plate forth or back and tilting it vertically or horizontally relative to the initial frontal view. This enhances calibration robustness by capturing variations in perspective. The angles between the views should be as large as possible, whereas keeping all markers on the calibration plate in focus. It is recommended that the markers of the calibration plate fill the whole region of interest required for the later measurement for each camera position, however it is of highly importance that the 3 disk-shaped markers located at the center of the plate are always visible by all cameras in all the robot positions acquired. Figure 5.7 shows the calibration plate used and the 4 views recorded for executing the robot geometric calibration.

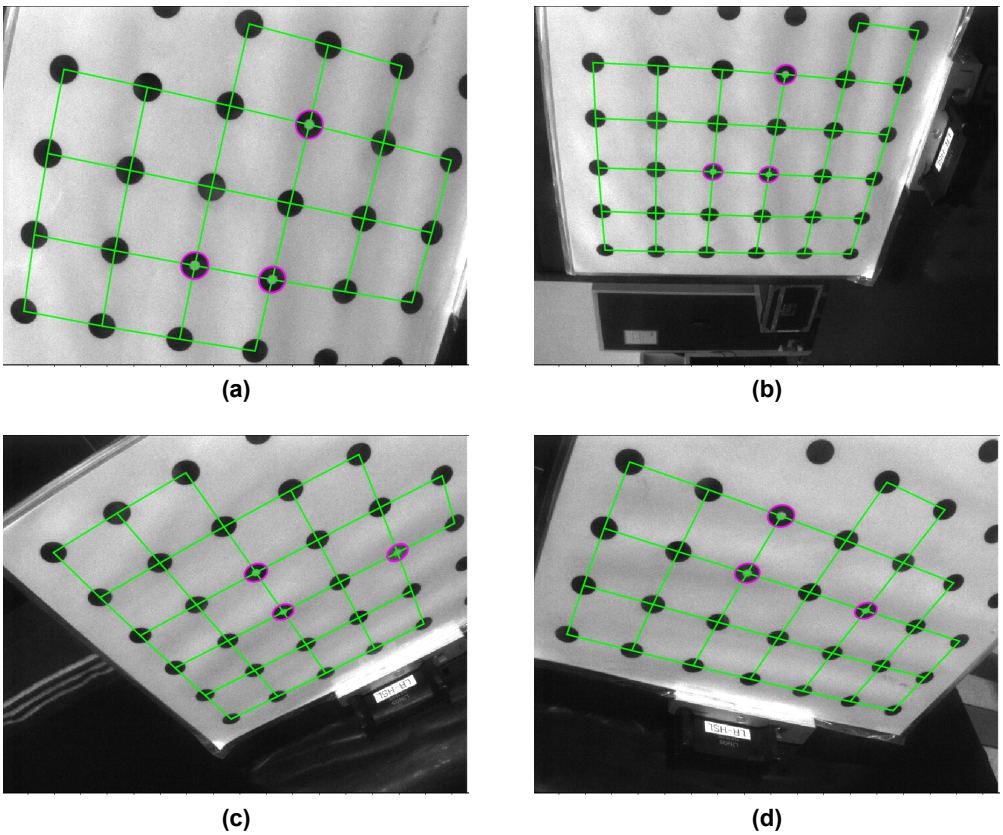

Figure 5.7: Calibration plate acquired images.

The system does not require calibration after repositioning of the CVV probe. In the case of moving the whole Robotic PIV system (CVV + robotic arm), such calibration must be re-done.

Volume Self-Calibration (VSC)

After the geometric calibration, the called *Volume Self-Calibration* has to be performed to correct and refine the initial calibration and reduce calibration errors to below 0.1 pixels [35]. A set of images where the robot only captures particle tracks (no models and avoiding reflections in all camera views) as shown in Figure 5.8 is recorded. The particle images are detected and triangulated in 3D space, yielding the preliminary 3D positions. The process involves back-projecting triangulated particles onto the different camera views and determining the differences between the back-projected particles and the original particle peaks. These differences indicate the degree of decalibration among the cameras. By averaging these differences across a large number of particles, disparity vectors are obtained. These vectors are then used to refine the initial image-to-world mapping functions obtained from the geometric calibration.

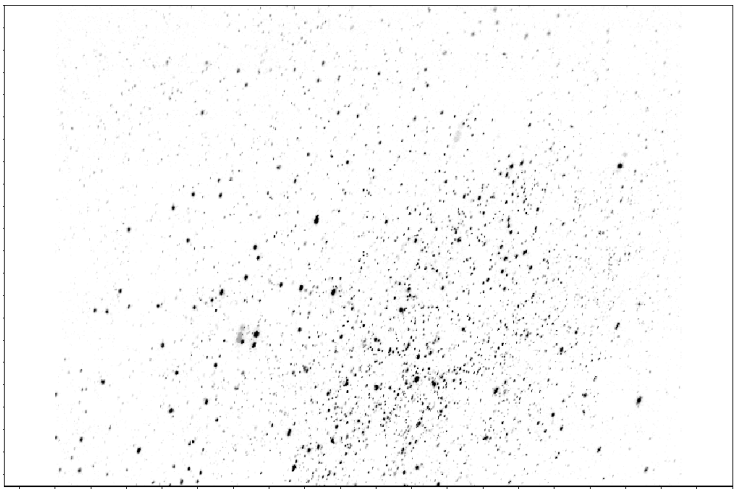

Figure 5.8: Volume Self-Calibration instantaneous image with recorded particles.

5.4. Data acquisition

For the side-view mirror and Formula 1 car models, the wind velocity is set at 12 m/s (following the study done by Saredi et al. in [93] with the same side-view mirror model). In the side-view mirror case, 8 robot positions are measured, whereas 10 views are acquired for the Formula 1 car. For both models, 5,000 images are recorded in sequence at each robot position in Double-Frame Shake-the-Box (DF-STB) mode with the following two time steps: $dt1$ = 100 μs and $dt2$ = 500 μs.

The propeller model is tested at a wind velocity of 5 m/s, as it is of use and interest to continue Cueto's study [94]. At the same time, the propeller rotational speed is set to 3600 RPM. The robot sequence consisted of 9 measurement volumes with 5,000 images recorded at each view in Time-Revolved Shake-the-Box (TR-STB).

The test matrix of all the experimental campaign is presented in Table 5.2.

Test number	Model	Velocity [m/s]	Propeller [rpm]	Acquisition	Images Acquired	Number of views
1	Side-view Mirror	15	-	Double frame	5,000 ($dt1$ = 100 μs) 5,000 ($dt2$ = 500 μs)	8
2	F1 Car	15	-	Double frame	5,000 ($dt1$ = 100 μs) 5,000 ($dt2$ = 500 μs)	10
3	Propeller	7	3600	Single frame	5,000	9

Table 5.2: Test matrix.

5.5. Data processing and reduction

As the aim of the experiment is to test and validate the proposed reflection treatment methodologies, the latter together with typical state-of-the-art image pre-processing approaches are applied to the acquired images. Hence, the raw images are pre-processed with four different techniques: minimum subtraction over the entire series of images, Butterworth time filter with a filter length of 9 images, spatial Fourier filter with a cut-on wavenumber κ_f = 0.1 px^{-1} and the spatial Fourier filter + mask. Then, the obtained images are processed with the particle tracking algorithm Shake-the-Box (see Section 2.1.1.2). Due to processing time limitations, only 500 images are processed for each method and test object. To remove outliers, velocity range and median filters are applied to the resulting data. The data obtained from Shake-the-Box shows the velocity of particles tracked over time as scattered data over the measurement domain. For the sake of simplicity and to facilitate interpretation of the flow field, a binning step is performed to transform the unstructured velocity information to a structured one.

The measurement volumes acquired for the side-view mirror are averaged in space and time within sphere-shaped voxels (or bins) of 15 mm diameter with a 75% overlap of neighbouring voxels, yielding a spacing of 3.75 mm between velocity vectors. The same binning is applied to the Formula 1 car case. For the propeller, the track data is interrogated within spherical bins with a diameter of 24 mm and 75% overlap (resulting in 6 mm of space between velocity vectors). Following Agüera et al. [98], a quadratic (2nd order) weighting function is employed on all cases to perform the velocity averaging inside each bin. Additionally, the set of measurement volumes are merged into a single dataset and the binning with the same parameters as for the individual views is applied.

6

Results and discussion

This chapter presents and discusses the results obtained from the experimental campaign, where three different geometries were tested: side-view mirror, Formula 1 car and propeller. The raw data acquired has been processed (considering the parameters from Section 5.5) with the aim to validate and check the performance of the proposed reflection treatment methodologies. Four image pre-processing techniques are compared: Minimum Subtraction over entire series (generates a background image with the minimum and then subtracts it from all images), Butterworth time filter (see Section 2.2.1.1), Spatial Fourier Filter and Spatial Fourier Filter + Mask[1]. Then, in the STB results, the 3D-based Particle Concentration Mask is also compared. The results are presented in separate sections, each corresponding to the individual test objects.

6.1. Side-view mirror case

For the side-view mirror geometry, eight different views were acquired in DF-STB mode with the robot. The sequence of measurement volumes is created such that it captures most of the flow field around the mirror, specially focusing on the wake. Figure 6.1 shows the individual measurement volumes (denoted in green) for the side-view mirror.

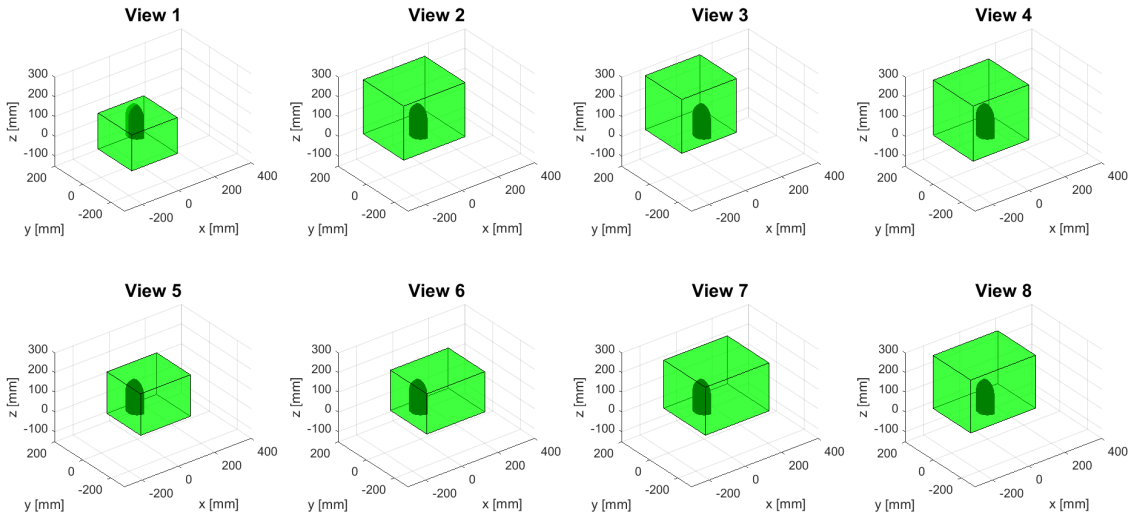

Figure 6.1: Set of measurement volumes acquired for the side-view mirror.

[1] Henceforth referred to as *SFF* and *SFFM*, respectively

For further details of what it is captured in each view, an instantaneous raw recording of each of the acquired views along with their pre-processed images are presented in Figure B.1 in Appendix B. All raw images show distinct and clear reflections that are easy to detect and identify by human eye. View 1 contains a large dense reflection caused by the floor (CVV probe too perpendicular to this surface) and the whole set of views presents a reflection due to the HFSB foam adhered to the surface of the mirror. The Minimum Subtraction fails to remove the reflections in all views. While the Butterworth filter effectively eliminates reflections in most views, it struggles with view 1, where the reflection is large and slightly unsteady, making it difficult to completely eliminate using this method. Regarding the proposed reflection treatment methodologies, the Spatial Fourier Filter does reduce the intensity of unwanted regions. However, as discussed in Chapter 4.1, it is unable to remove the high-wavenumber components present in reflections. Consequently, residual reflection components are still visible in the Spatial Fourier Filtered images, especially in view 1. The SFFM images demonstrate that the method can robustly remove the reflection region from all views by creating a mask around them. Hence, blank empty regions can be observed in these images. Recall that this technique presents the drawback of removing real particles that pass over the spurious region that is being masked. Nevertheless, it shows great potential in terms of adaptive masking, and to check this capability, a set of four consecutive recordings for views 2 and 3 are shown in Figures 6.2 and 6.3, pre-processed with the Butterworth and Spatial Fourier Filter + Mask approaches.

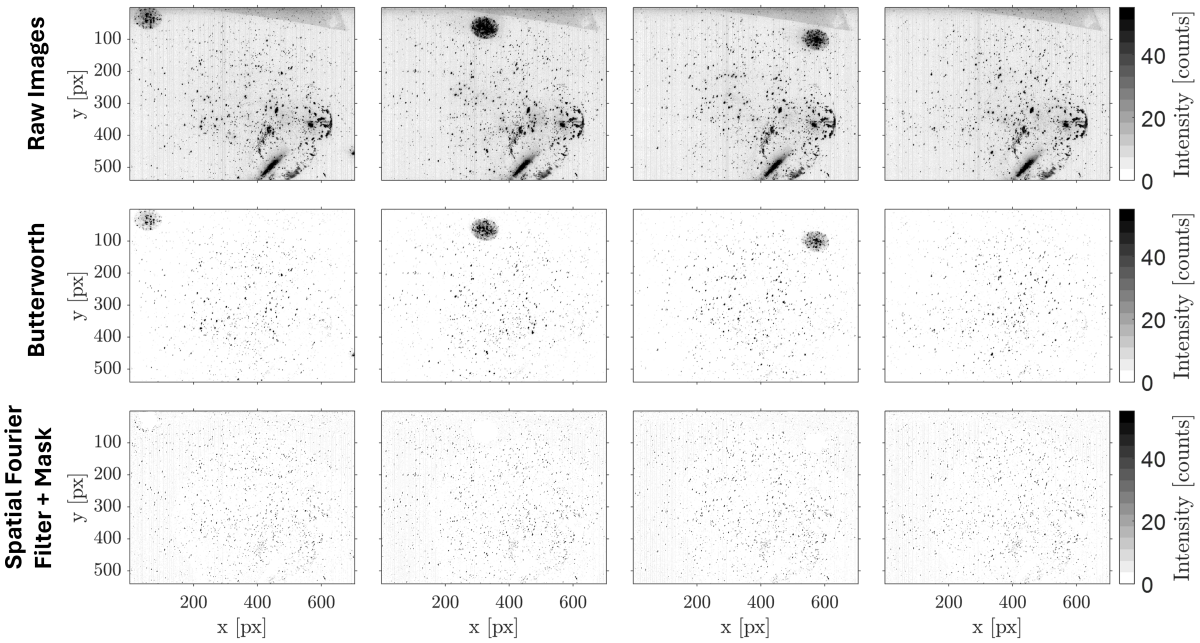

Figure 6.2: Set of consecutive pre-processed images (from left to right) with no filter, Butterworth time filter and Spatial Fourier Filter + Mask (View 2).

Some views suffered the presence of large foam blobs caused by the malfunctioning of a few nozzles, and the images above show two examples of this for the case of view 2 and view 3. These foam blobs appear as big areas of high intensity, easily distinguishable from particles and their surroundings. They pose a significant challenge to particle tracking, as their appearance is random and can blind particles being tracked, thus interrupting the tracking process. For this reason, it is essential to remove these sporadic artifacts when they make their appearance. As they appear and disappear at a similar rate of the particle tracks, the Butterworth time filter is unable to mitigate the blobs (as shown in the top row in Figures 6.2 and 6.3). However, the

SFFM clearly is able to identify these regions in each acquired image and, then eliminate them by masking them. This empty space in the image can be filled in thanks to the multiple-view capability of Robotic PIV. By measuring views that overlap, regions that are blank in one view can contain data in another view, thus making it possible to fill in the gaps when merging all the measurement volumes.

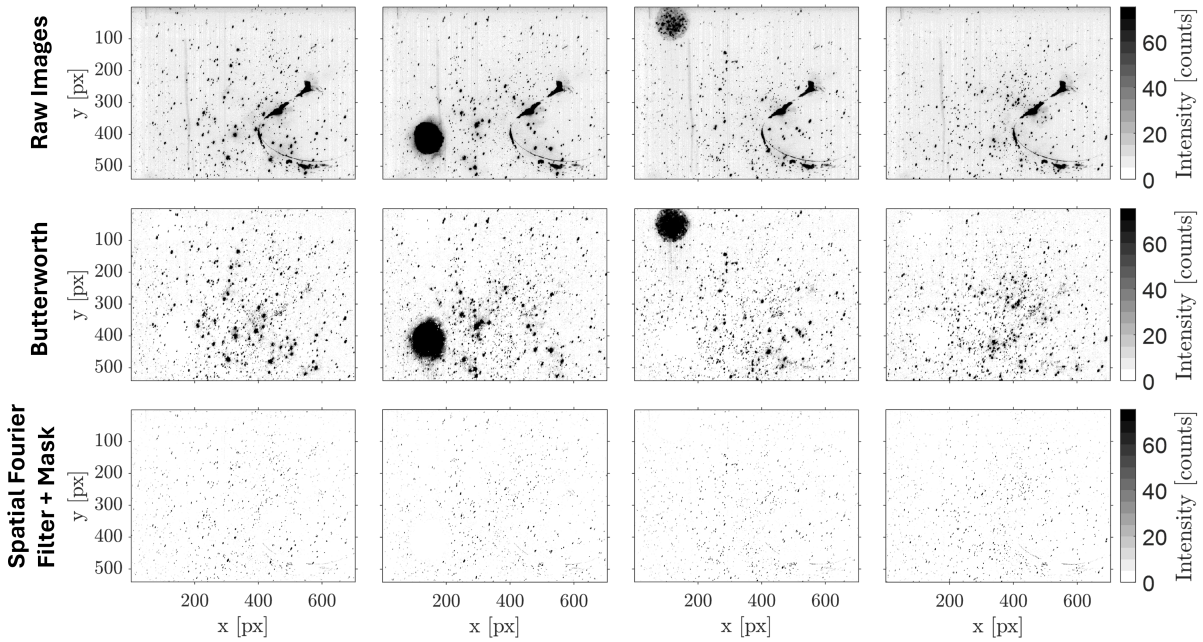

Figure 6.3: Set of consecutive pre-processed images (from left to right) with no filter, Butterworth time filter and Spatial Fourier Filter + Mask (View 3).

Figure 6.4 presents the number of tracked particles by the Shake-the-Box algorithm for each recording and view, and comparing the different reflection treatment methodologies considered.

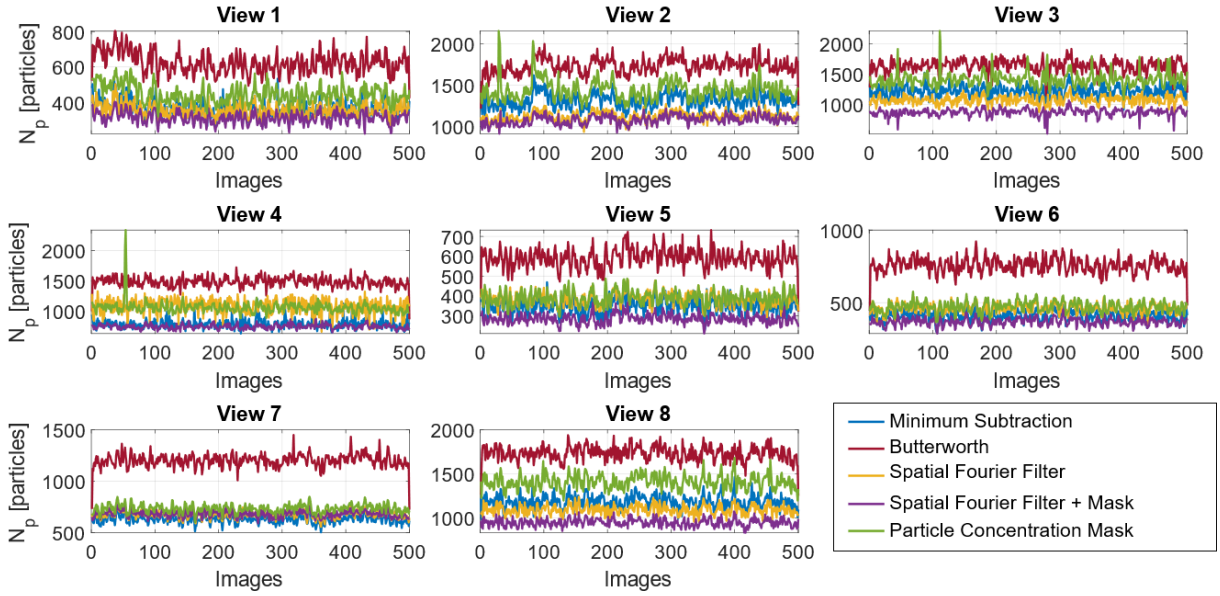

Figure 6.4: Number of tracked particles per recording for each view (side-view mirror case).

Two key observations can be made from the plots: the Butterworth time filter yields the highest number of tracked particles, while the SFFM method provides the least. This result is logical because the side-view mirror is a steady reflections case, where the Butterworth is able to robustly remove the reflections regions from the images. Conversely, in the case of the Fourier filter + mask, as a mask is being applied, reflections and particles are being eliminated, which can interrupt the detection and tracking of the latter. This explains why this method tracks the fewest number of particles compared to the others, with a common trend of SFFM approach tracking approximately half the particles compared to the Butterworth. Regarding the Minimum Subtraction and Particle Concentration Mask, both show similar tracked particles in most of the views, which can be translated into that the second fails to mask the reflection.

6.1.1. Individual views

Let us check the effect of each method on a few individual measurement volumes. Figure 6.5 shows the raw and pre-processed images of View 1.

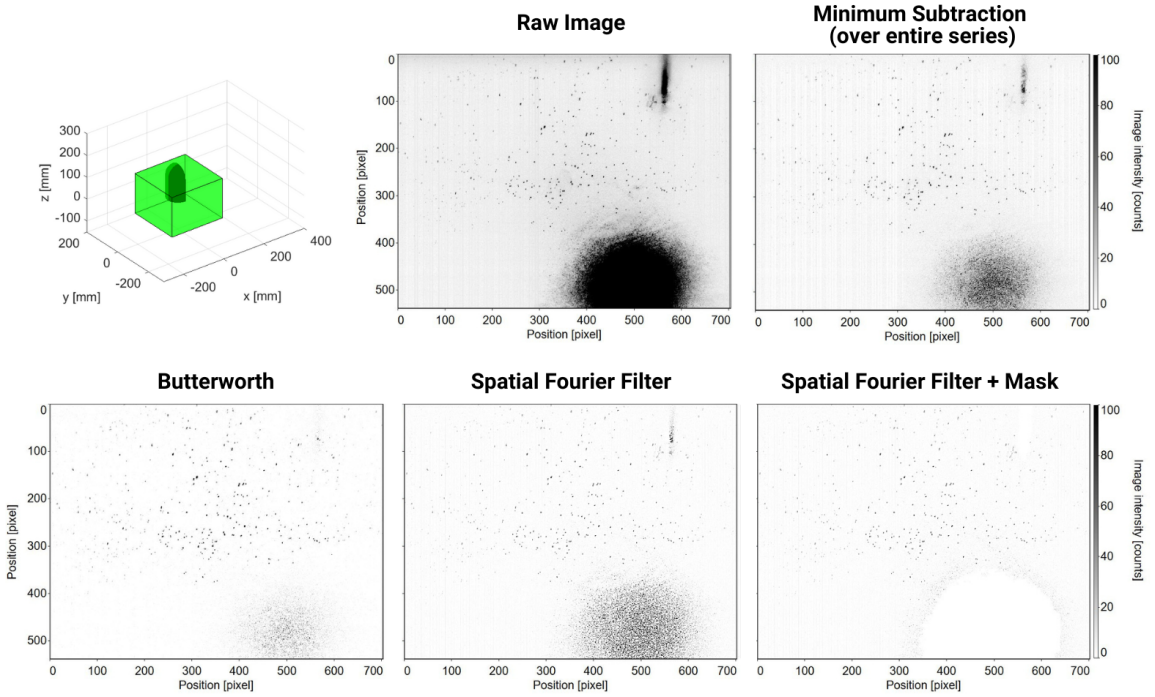

Figure 6.5: Side-view mirror acquired and pre-processed images of View 1.

View 1 has the presence of two main reflections: bottom large one that originates on the floor and top smaller one coming from the foam adhered on the mirror surface. Neither of these are removed by the Minimum Subtraction or the Spatial Fourier Filter, moreover in the regions of the floor reflection there are still a big amount of particles that will surely slow down the Shake-the-Box process. The Butterworth time filter attenuates both reflections, but still does not removes them completely. And the Spatial Fourier Filter + Mask successfully creates a mask of the two spurious regions and, thus removes their appearance. This reflection appears in the results shown in Figure 6.6 for View 1. All methodologies present regions that correspond to the reflection (some approaches in smaller size), except for the SFFM which successfully removes it.

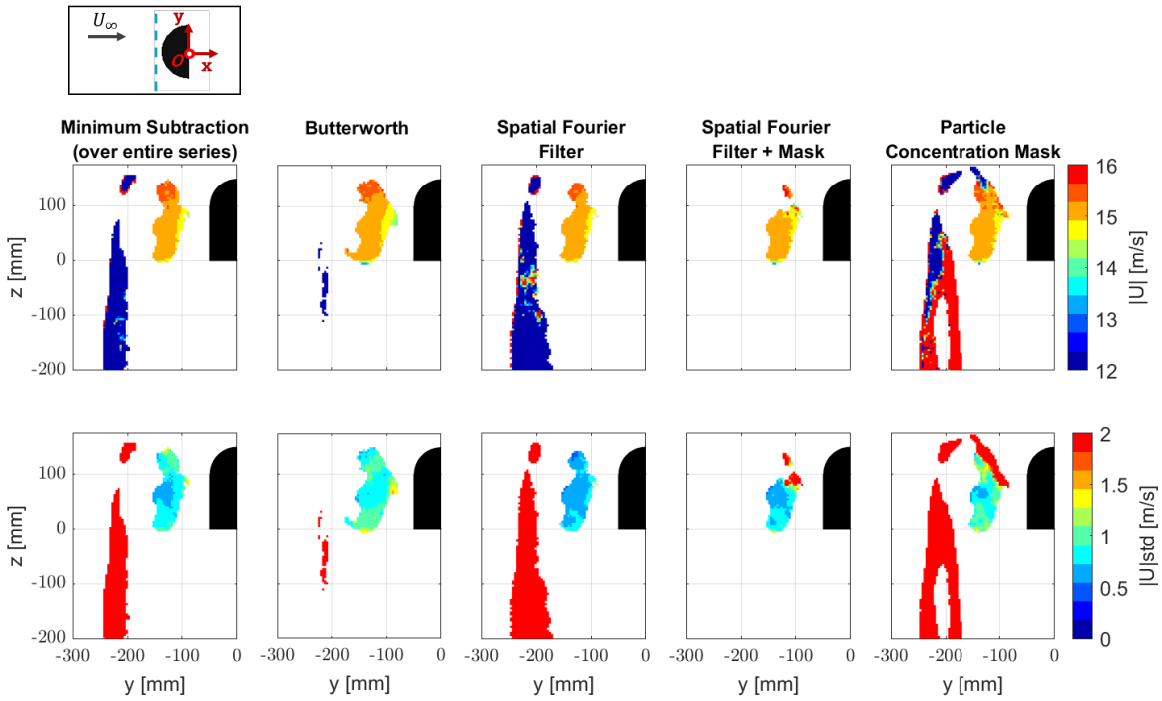

Figure 6.6: Side-view mirror binning results on the YZ plane at x = -50 mm (View 1).

Having discussed the results of View 1 with a clear large reflection, let us focus on measurement volume 8 that contains a smaller reflection, almost not perceived in the Butterworth and SFFM cases, as shown in Figure 6.7. This measurement volume mainly focuses on the wake of the

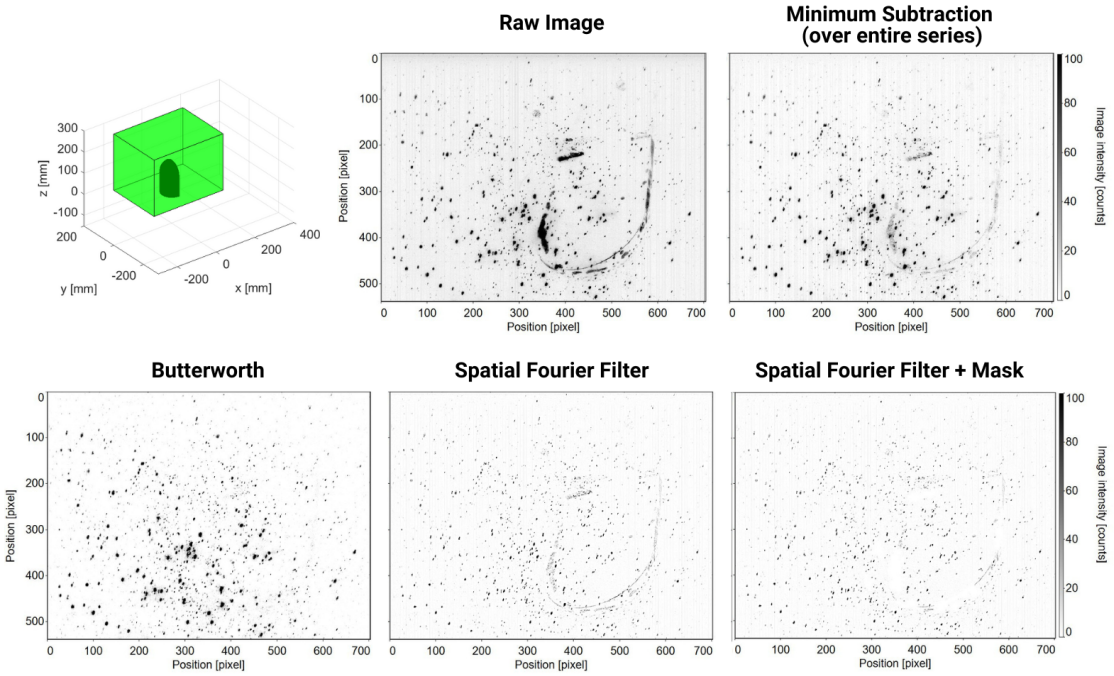

Figure 6.7: Side-view mirror acquired and pre-processed images of View 8.

mirror, capturing the back and a portion of the top of the mirror. In the acquired images, there are reflections that originate from the foam stuck on the mirror surface and thinner ones caused by the edge of the mirror. As with View 1, the Minimum Subtraction and Spatial Fourier Filter are not able to remove fully these regions, whereas Butterworth and SFFM are. The latter presents blank

areas on the edges of the mirror and where the foam was captured in the image. However, the reflection on the edge of the mirror is not masked in all images equally since this region can be too small in size, with its resulting area in px^2 being too similar to some of the biggest particles. Therefore, in the step of the SFFM where only the biggest artifacts are kept, the reflections from thin edges can be excluded.

Figures 6.8 and 6.9 show the XZ and YZ planes, respectively, of the mean velocity and mean velocity standard deviation for View 8.

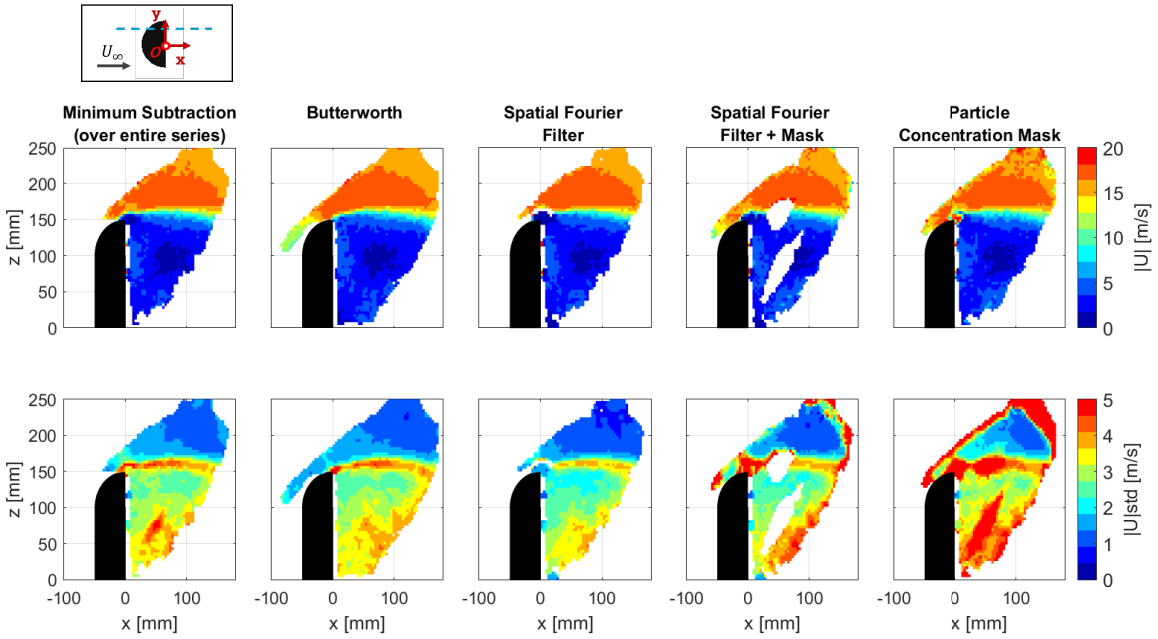

Figure 6.8: Side-view mirror binning results on the XZ plane at y = 0 mm (View 8).

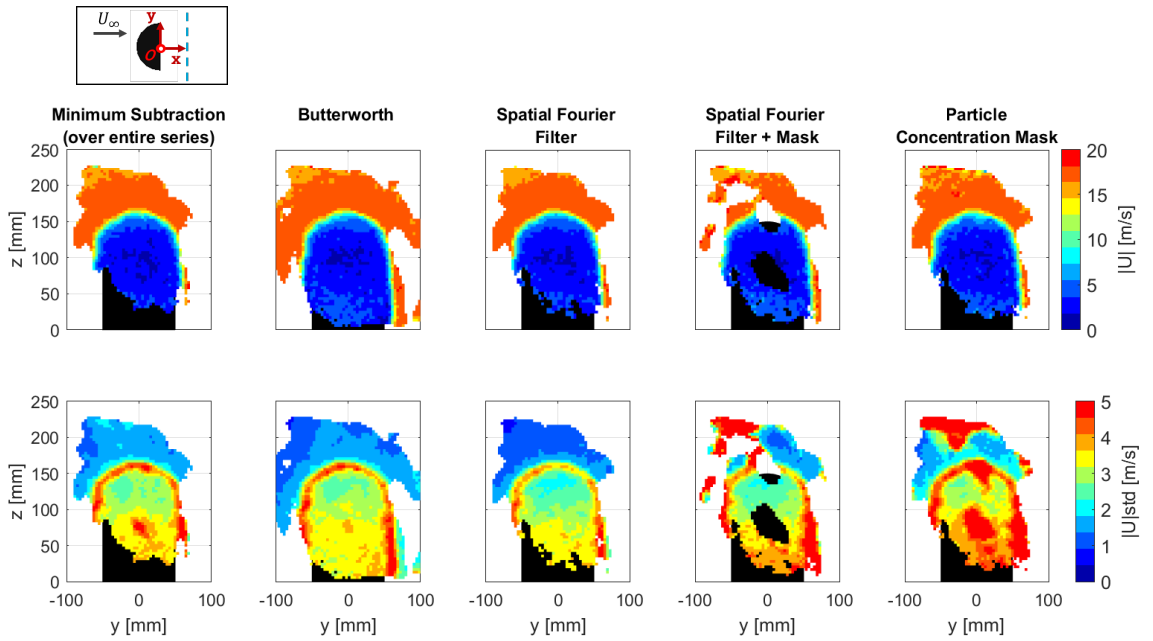

Figure 6.9: Side-view mirror binning results on the YZ plane at x = 50 mm (View 8).

All the methods, except SFFM, show similar results in the velocity field. Minimum Subtraction, Spatial Fourier Filter and Particle Concentration Mask approaches fail to remove the regions due to the foam stuck on the mirror surface, with the latter approach yielding a higher standard deviation compared to the other methods. The SFFM result shows a particular characteristic: there are voids within the data (resembling the appearance of Emmental cheese), consequence of masking regions in the images processed with Shake-the-Box – see Figure 6.10. The algorithm cannot find anything in this area, which translates into in missing data in the resulting flow field. This can be seen as a clear drawback (since having gaps in the data is detrimental to a good understanding of the flow field), however thanks to the advantageous characteristic that Robotic PIV offers of being able to easily acquire measurement volumes from different views, these empty regions can potentially be filled in by another measured view.

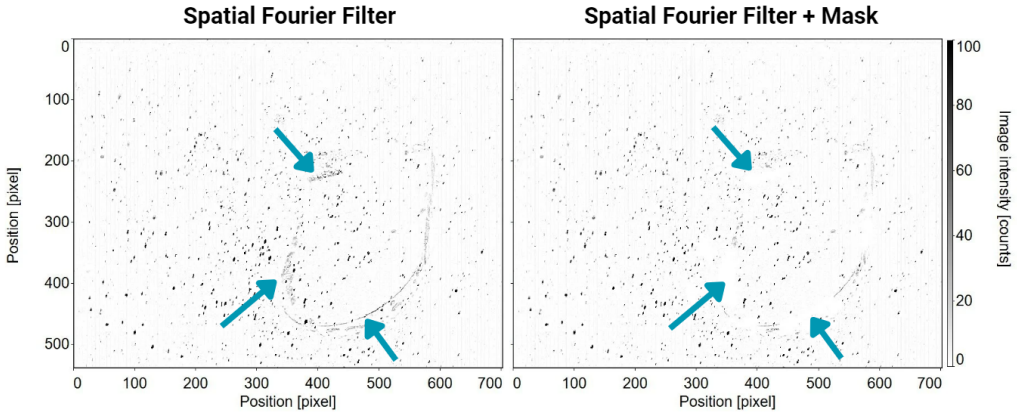

Figure 6.10: Close-up view of Spatial Fourier Filter and Spatial Fourier Filter + Mask methods in View 8.

6.1.2. Complete measurement domain

Once all the individual views are processed, the complete measurement volume around the side-view mirror can be obtained. DaVis software offers the option called "*Particle Track Stitching*" that merges the set of tracks of a multi-set (more than one view) to one joint set. After the merged set is obtained, the binning operation with the same parameters as used for the individual views can be performed.

Figures 6.11, 6.12 and 6.13 show the XZ, YZ and XY planes, respectively, of the particle concentration, mean velocity and mean velocity standard deviation for joint dataset of measurement volumes. The Particle Concentration Mask approach performs poorly, failing to remove reflection regions and resulting in an incorrect flow field. The Minimum Subtraction, Butterworth, Spatial Fourier Filter, and Spatial Fourier Filter + Mask present a similar mean velocity distribution. However, there is a notable difference between the standard deviation of these methods, with the Spatial Fourier Filter and Spatial Fourier Filter + Mask exhibiting a lower standard deviation compared to the other three methods. Recall the results of View 8 for the SFFM that showed holes in the data in the three planes displayed. This was due to the masking step on the images that yielded no data after processing with Shake-the-Box. After performing the merge of data, it is proven that this gaps are filled in by the data of other views since there are no empty spaces in the flow field for the SFFM in Figures 6.11, 6.12 and 6.13 (fourth column). This is thanks to having acquired several views whose volumes in space overlap and, thus allow to compensate the loss of data due to masking.

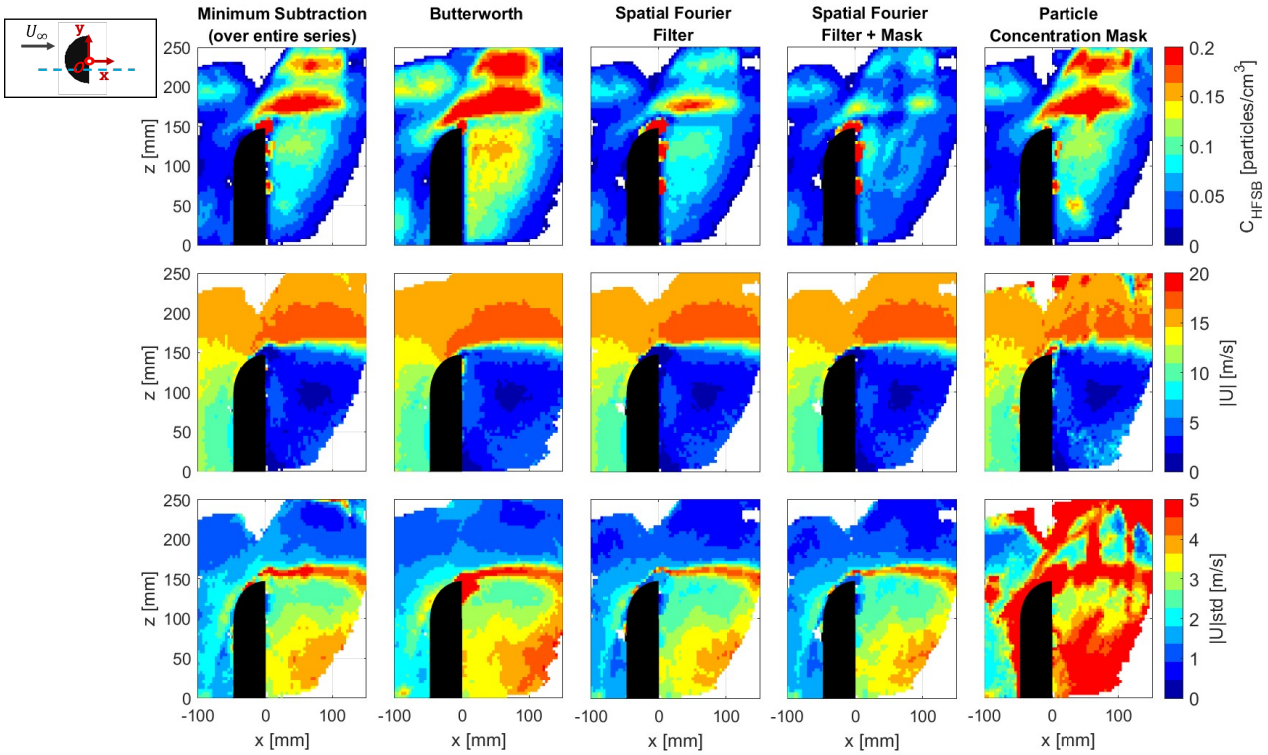

Figure 6.11: Side-view mirror binning results on the XZ plane at y = -10 mm (all views).

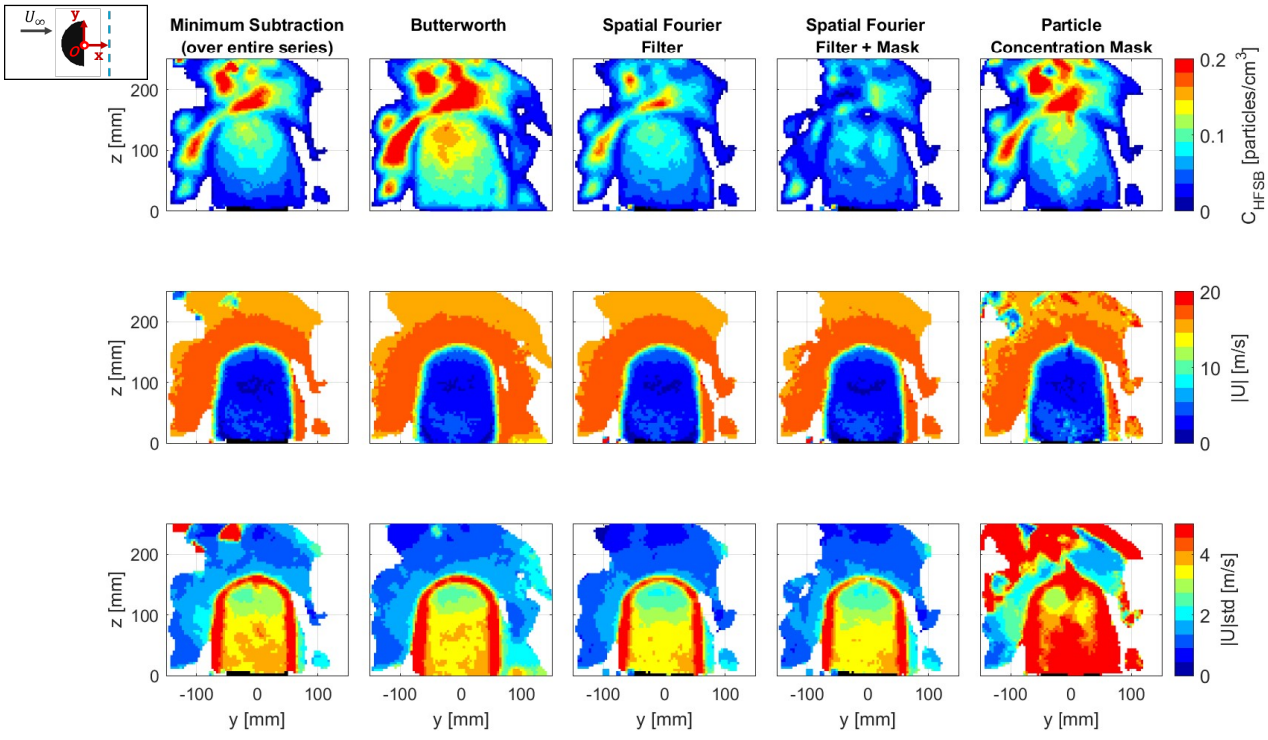

Figure 6.12: Side-view mirror binning results on the YZ plane at x = 50 mm (all views).

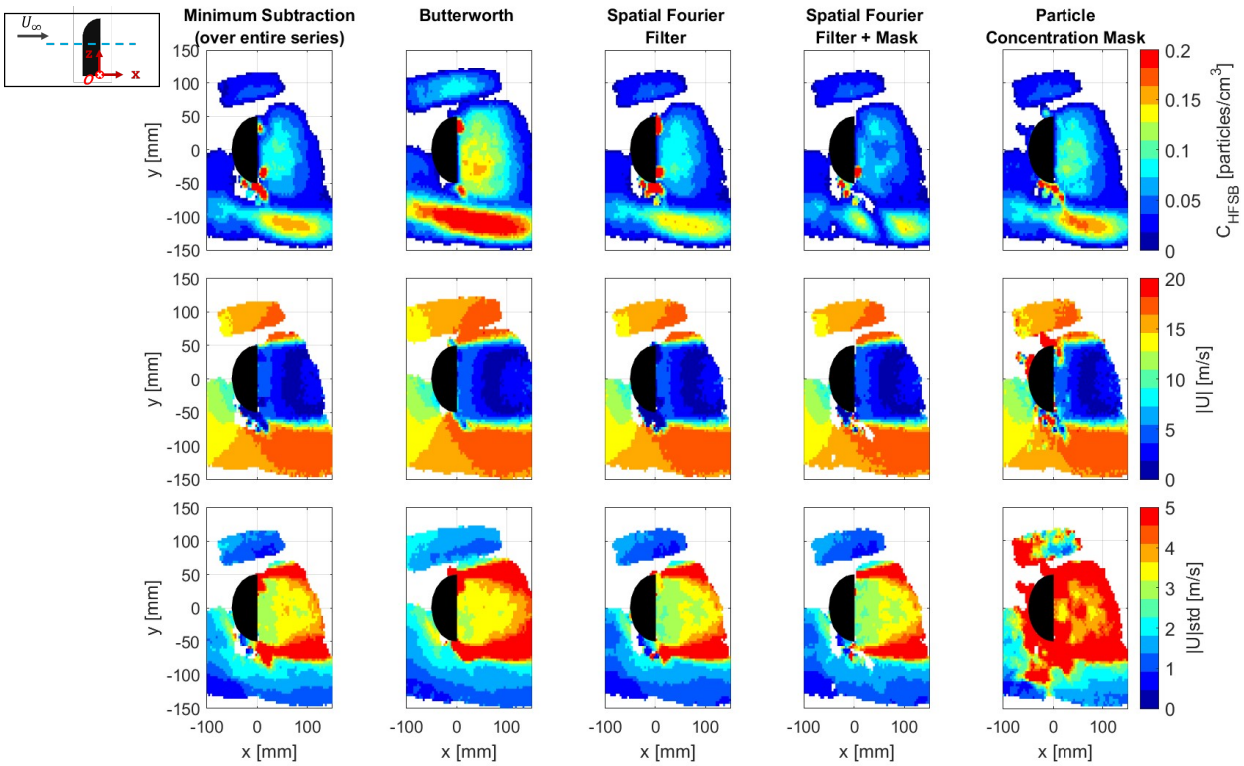

Figure 6.13: Side-view mirror binning results on the XY plane at z = 85 mm (all views).

It is worth noticing that the Field of View (FOV) for the Spatial Fourier Filter and Spatial Fourier Filter + Mask is smaller in all planes shown compared to the Butterworth. Both upstream and downstream regions show a more complete velocity field for the Butterworth, yielding a longer wake behind the mirror. This is an important characteristic since it allows to retrieve as much information as possible with the measurement volumes acquired.

When inspecting the processed images, in the case of SFF and SFFM the particles close to the edges of the image have a decaying intensity with respect the ones in the center; for instance, inhomogeneous laser illumination can be a cause of this effect. However, this is not observed for the Butterworth, which might explain its resulting larger FOV as it allows the particle tracking algorithm to detect more easily the particles further from the center of the image. This suggests that Butterworth operation performs an intensity normalization after having applied the time filter on the images to counteract nonuniform intensity over the image. One way to avoid this the decaying particle intensity effect in the proposed image pre-processing approaches is to add an additional step of intensity normalization on the images.

Therefore, to check that the FOV loss is due to inhomogeneous image intensity, an intensity normalization (over 100 pixel) is applied to the SFFM images to make the intensity uniform across the image. The results of this operation can be observed in Figures 6.14, 6.16 and 6.15.

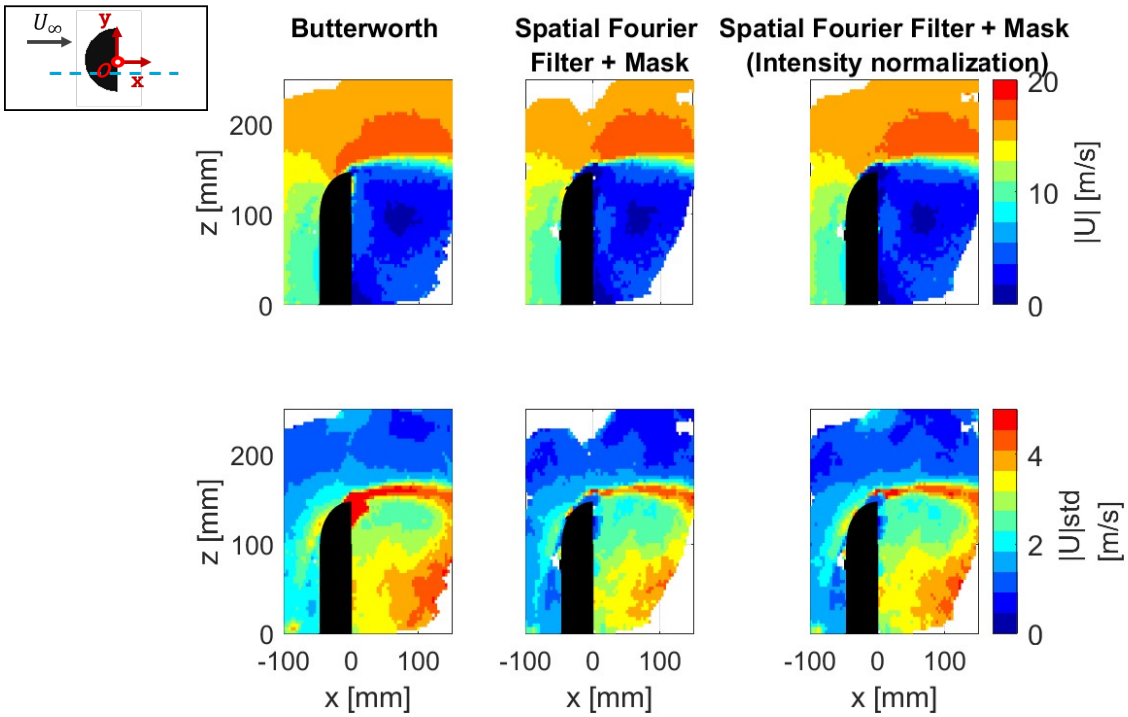

Figure 6.14: Side-view mirror Butterworth, SFFM and SFFM with intensity normalization results on the XZ plane at y = -10 mm (all views).

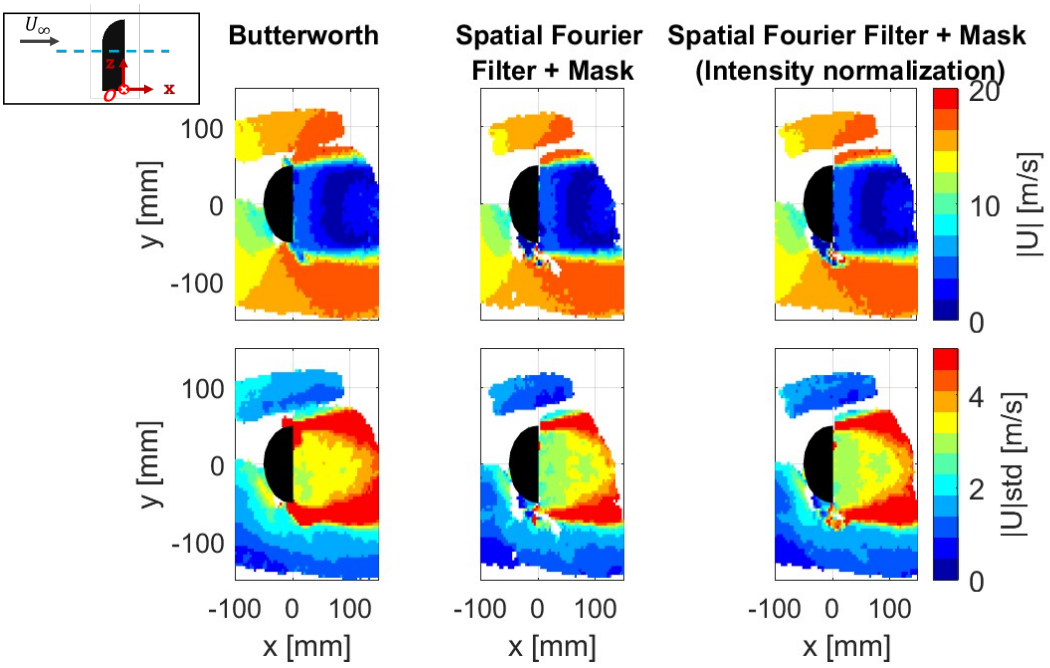

Figure 6.15: Side-view mirror Butterworth, SFFM and SFFM with intensity normalization results on the XY plane at z = 85 mm (all views).

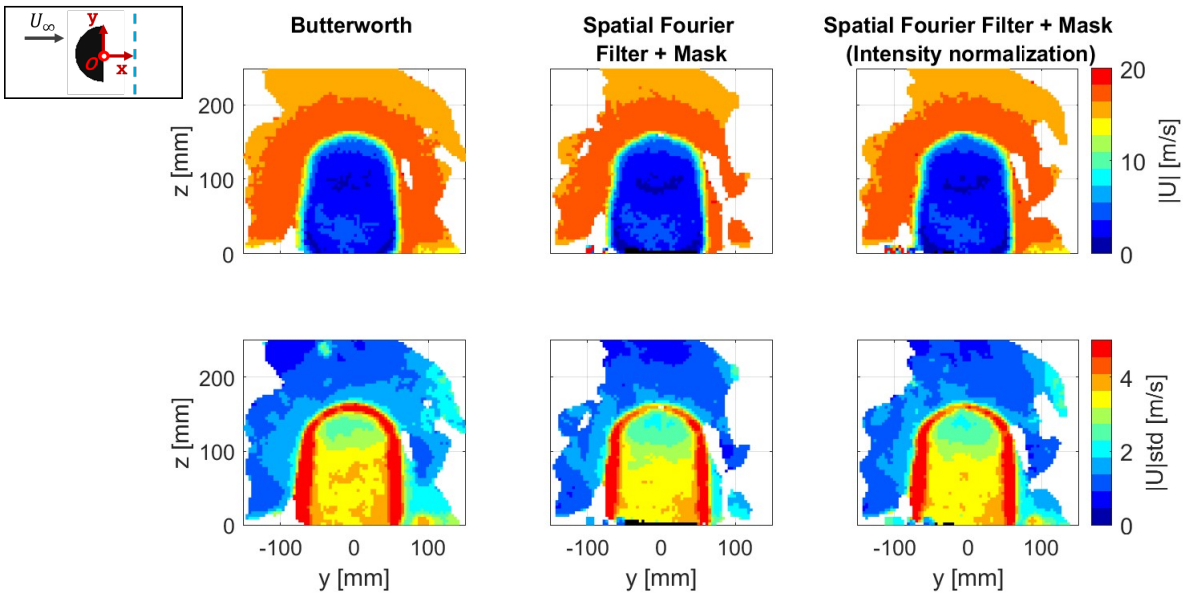

Figure 6.16: Side-view mirror Butterworth, SFFM and SFFM with intensity normalization results on the YZ plane at x = 50 mm (all views).

After applying intensity normalization to the SFFM images, the resulting flow field expands in size in all displayed planes. Although the field of view remains smaller than that of the Butterworth filter, it is noticeably larger than the not normalized SFFM result. Furthermore, the standard deviation shows lower values compared to the Butterworth filter. This indicates that SFFM with intensity normalization is an effective reflection treatment methodology, producing similar results to those of the Butterworth time filter.

6.1.3. Further discussion on Butterworth - SFFM

Among all the reflection treatment methodologies tested, the Butterworth time filter and the spatial Fourier filter and masking approach show the better performance in removing the spurious regions of reflections. Therefore, in order to check their results with when processing a higher number of images, the 5,000 images acquired are processed for both the methodologies. The results of processing these images with STB are presented in Figures 6.17 and 6.18, by showing the binned velocity field on XZ and YZ planes, respectively. The results are given for three single views as well as the merged dataset. For the latter, the standard deviation of the velocity magnitude is also given in the figure.

Overall, the resulting velocity field of the masked case does resemble the Butterworth's with a similar mean velocity in the single views and in the merged case, with almost no blank spaces appearing in the single views. The presence of no empty spaces is significantly important to notice since these appeared when only processing 500 images. This is thanks to the amount of data when processing the set of 5,000 images. Therefore, it is key to take into account this feature when implementing the spatial Fourier filter and mask approach.

The Butterworth case shows a small region of outliers at the front of the side-view mirror in View 4. This is due to foam attached to the surface of the mirror that the time filter is not able to fully remove because of its fluctuating behaviour. Conversely, the spatial Fourier filter and mask approach is able to completely remove this region, avoiding its appearance in the processed results. This blank space in the 2D images is translated into an empty gap in the velocity field along the line of sight of the velocimeter (see Figure 6.18, View 4). However, this void is almost entirely filled in the merged dataset velocity field, thanks to the presence of other

views that contain information within this region of the domain. In this case, View 8 is an example of measurement volume that has data where there View 4 presents a gap (top left of the mirror).

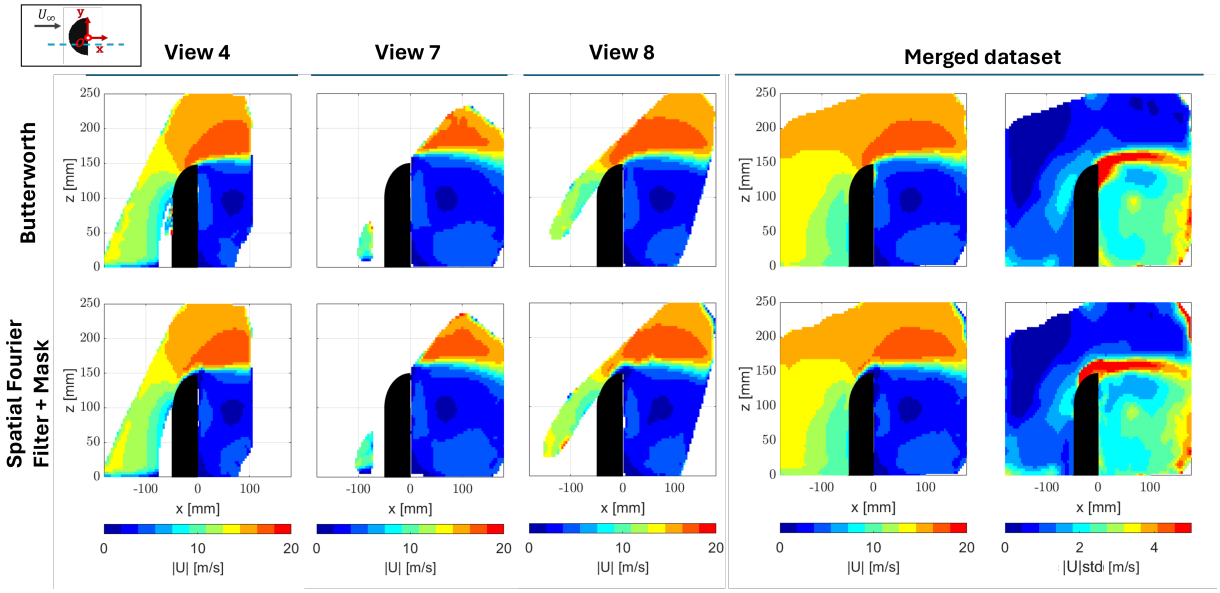

Figure 6.17: Side-view mirror binned results on XZ plane at y = 0 mm for single and merged views (Butterworth and Spatial Fourier Filter and Mask).

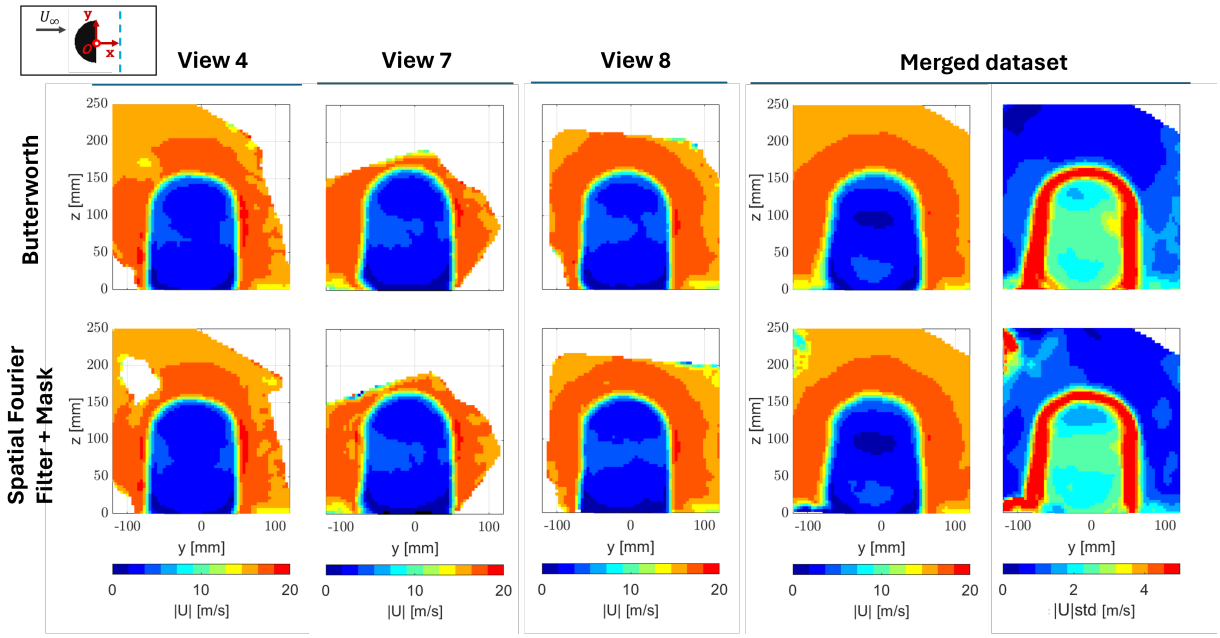

Figure 6.18: Side-view mirror binned results on YZ plane at x = 30 mm for single and merged views (Butterworth and Spatial Fourier Filter and Mask).

Comparing the merged datasets of both methods, the mean velocity is in excellent agreement, but the standard deviation presents slightly different features, particularly in the shear layer in Figure 6.17. The Butterworth shows high standard deviation in the wake close to the mirror, while the masking method exhibits high standard deviation on top of the mirror. This phenomena are due to the local effect from the presence of residual reflection not being fully removed. Despite of this, the proposed spatial Fourier filtering and masking method performs as well as the Butterworth time filter in the case of stationary objects.

6.2. Formula 1 car case

In the case of the Formula 1 car, ten views were measured in DF-STB mode with the robot. Figure 6.19 shows the individual measurement volumes (denoted in green) obtained for the F1 car.

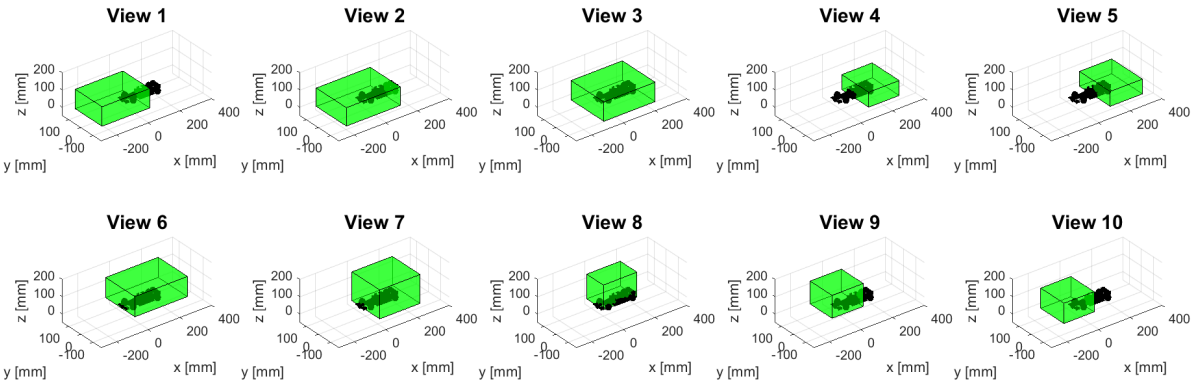

Figure 6.19: Set of measurement volumes acquired for the Formula 1 car.

Views 1, 2, 9 and 10 focus on the front of the car, views 3 and 8 on the main car geometry and views 4, 5, 6 and 7 capture mostly the rear. This can be observed on the raw images displayed in Figure C.1 in Appendix C. In all views, reflections originate from the Formula 1 car model and objects on the floor (e.g. screws). In the case of the side-view mirror, reflections appears more concentrated in certain regions, whereas for the F1 car, reflections are slightly more distributed, adopting the shape of the geometry. This characteristic is important to understand the resulting images processed with the Spatial Fourier Filter + mask methodology. For instance, View 1 shows a reflection from the front wing, front left tire and suspensions and also from a screw on the floor. As there are these many details, where some of them are smaller in size, the SFFM approach finds difficulty to identify and mask all the artifacts properly, which results in some of them not being mask in some of the views. In this case, reflections are mainly steady, which benefits the implementation of the Butterworth time filter.

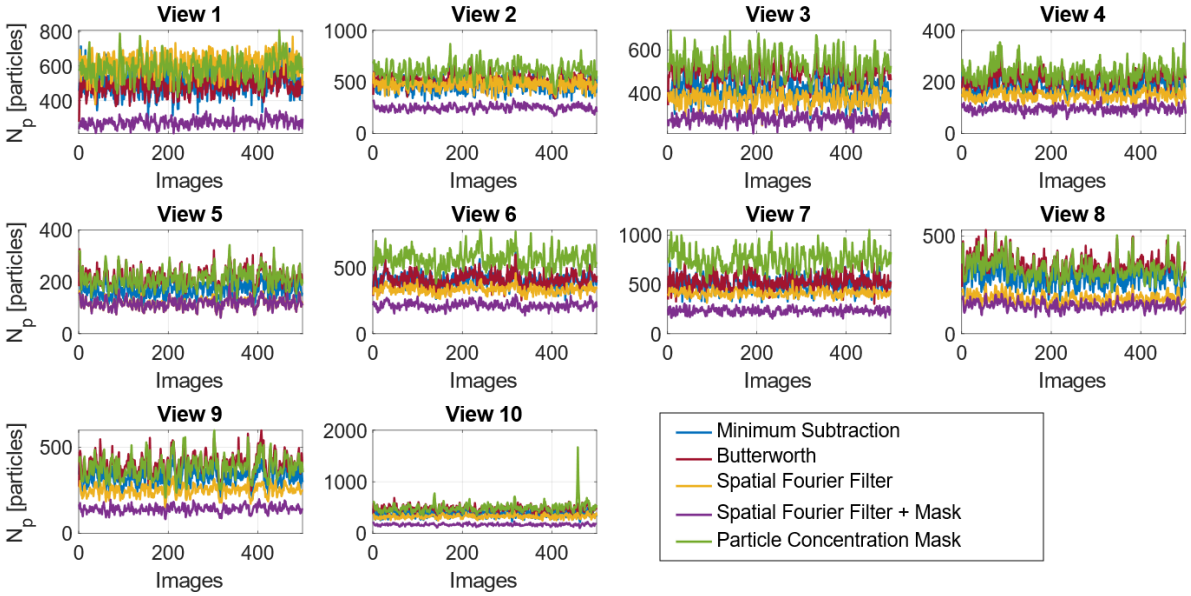

Figure 6.20: Number of particles per recording for each view (Formula 1 car case).

The number of particles tracked per recording is plotted for each view and reflection treatment method in Figure 6.20. Similarly to the side-view mirror case, the Spatial Fourier Filter + Mask approach presents the fewest number of particles tracked over the set of images, which is logical given that a region is being removed when applying the mask. Regarding the rest of methodologies, they approximately show similar results in all views, which does not allow to draw clear conclusions.

6.2.1. Individual view

All measurement volumes present similar reflection characteristics with a steady nature, therefore in this case only the results of a single individual view is analysed in this section. View 6 is selected to take a closer look as it is examples of volume measured at the rear of the car, yielding the velocity field at the wake. Figure 6.21 shows the raw and processed images for View 6. The images present a reflection that emanates from the back of the car, which are attenuated but not completely removed by the Minimum Subtraction and the Spatial Fourier Filter. The images resulting from the latter approach still contain residual artifacts from the reflection that the Fourier filter is not able to eliminate. However, when applying the extra step of masking with the Spatial Fourier Filter + Mask, these artifacts are almost fully mitigated. The Butterworth filter proves to be the most effective option (due to the reflection being steady), completely removing the reflection content from the image while retaining only the particle tracks.

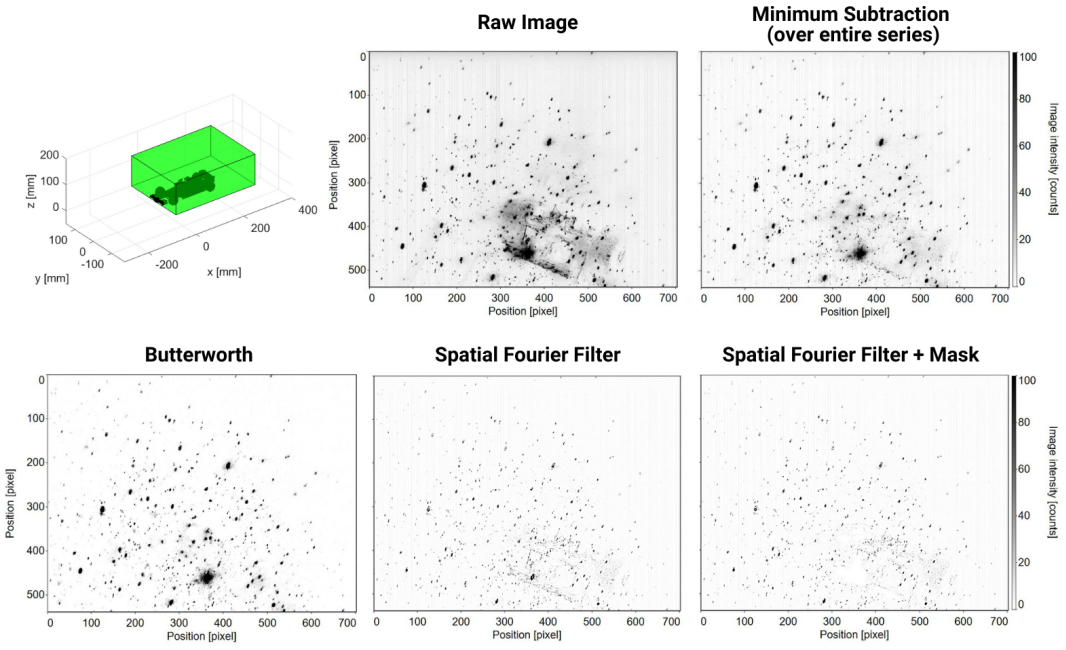

Figure 6.21: Formula 1 car acquired and pre-processed images of View 6.

The binning results of View 6 for the XZ, YZ and XY planes are shown in Figures 6.22, 6.23 and 6.24, respectively. The Particle Concentration Mask approach fails to remove the regions of reflections, resulting in an erroneous velocity field and a much higher standard deviation compared to the other methods. Regarding the other techniques, they all show similar results with a similar velocity field: low velocity zone at the back of the car due to the presence of the rear wing. Although SFFM shows a similar result, it presents a blank space near the rear wing, which is due to masking reflection caused by this part of the car.

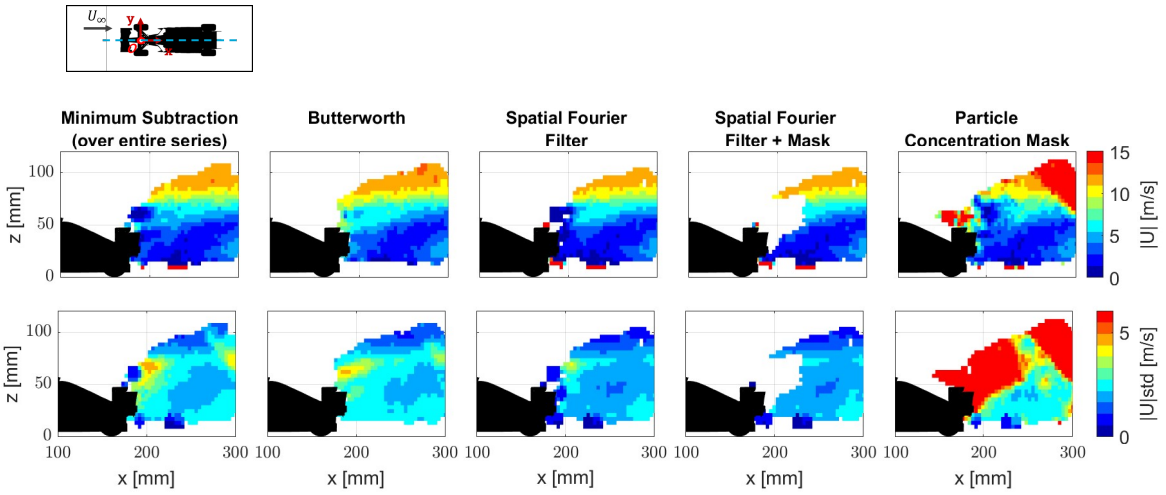

Figure 6.22: Formula 1 car binning results on the XZ plane at y = 0 mm (View 6).

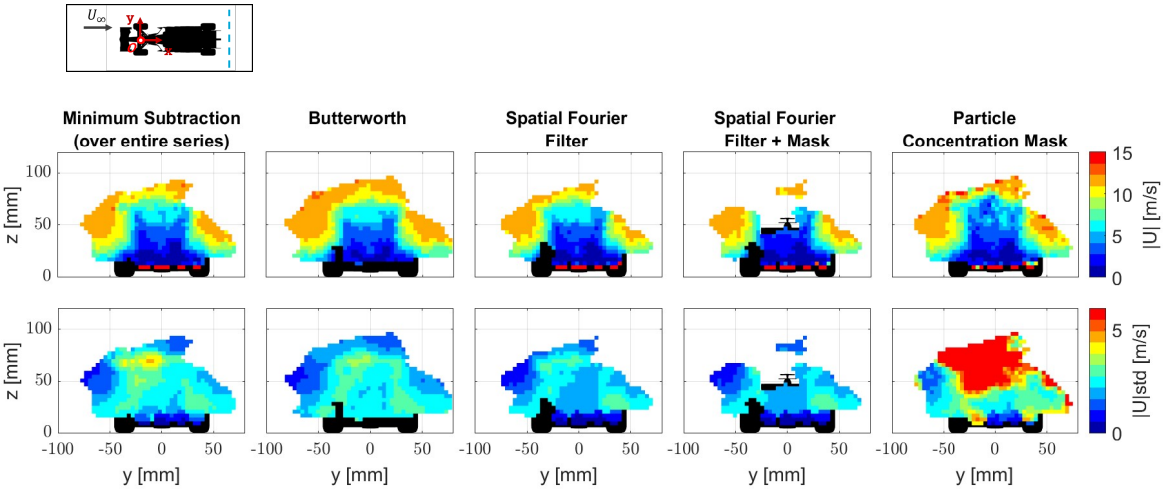

Figure 6.23: Formula 1 car binning results on the YZ plane at x = 220 mm (View 6).

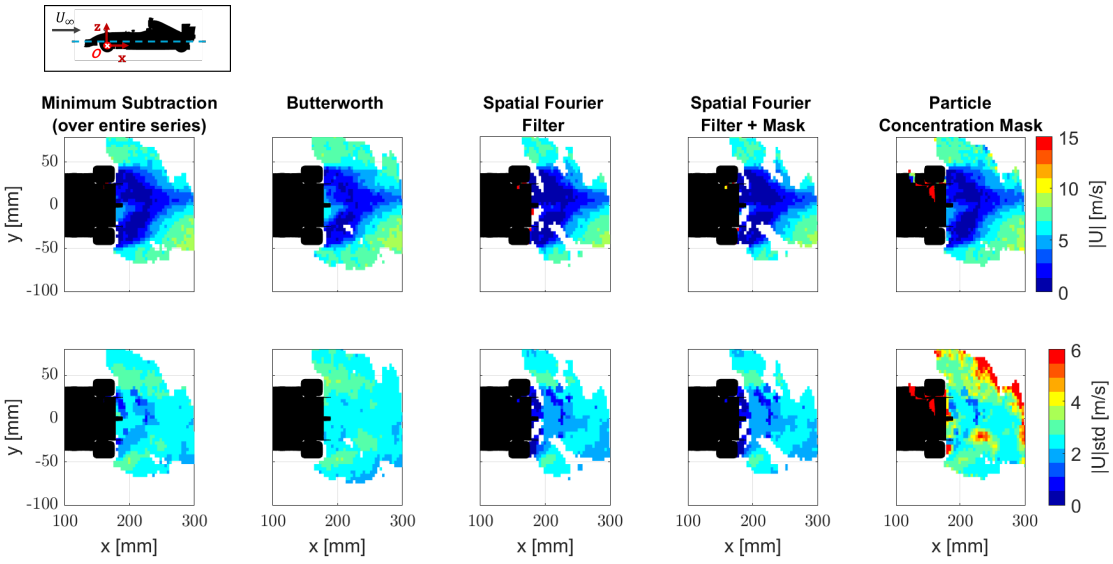

Figure 6.24: Formula 1 car binning results on the XY plane at z = 20 mm (View 6).

6.2.2. Complete measurement domain

After processing all the measurement volumes separately, the entire resulting flow field can be obtained by merging all the views. Then the binning operation with the same parameters as with the individual views is executed. The particle concentration, mean velocity and mean velocity standard deviation of the merged dataset are depicted in Figures 6.25, 6.27 and 6.26 for the XZ, YZ (rear), YZ (front) and XY planes, respectively.

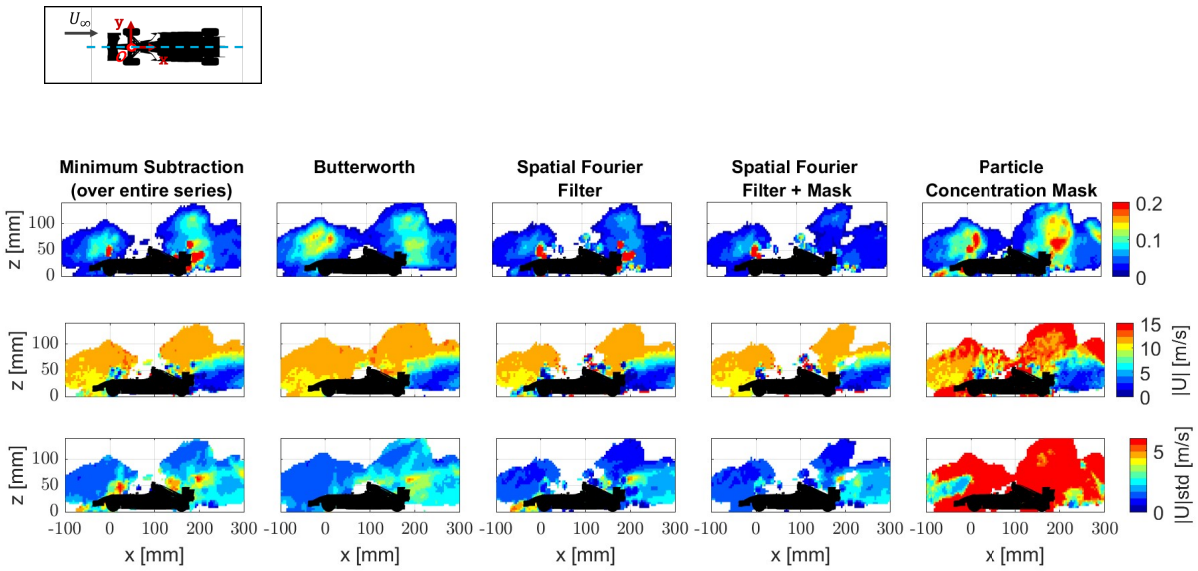

Figure 6.25: Formula 1 car binning results on the XZ plane at y = 0 mm.

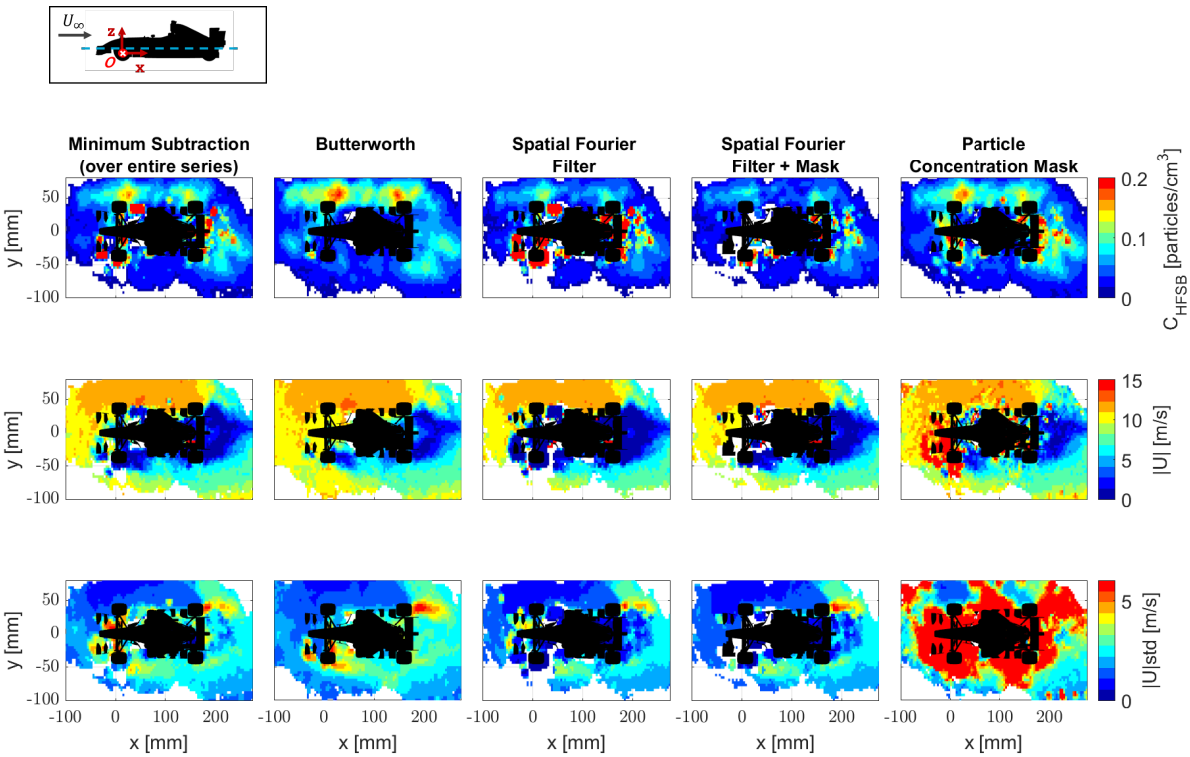

Figure 6.26: Formula 1 car binning results on the XY plane at z = 20 mm.

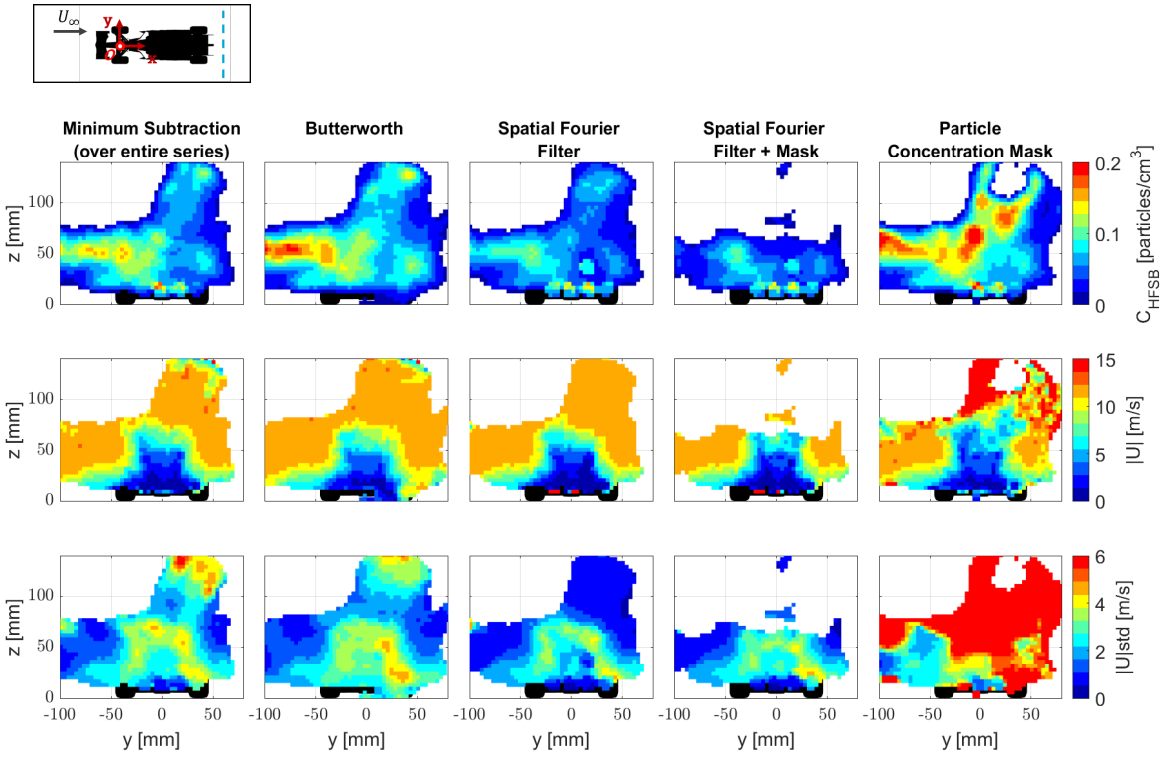

Figure 6.27: Formula 1 car binning results on the YZ plane at x = 220 mm.

Looking at the results from the Particle Concentration Mask, one can see that this method does not remove the reflection, but more importantly, the resulting flow field deviates from the expected one such that it leads to misunderstanding. The Butterworth case exhibits the most complete flow field with the largest field of view and some of a Formula 1 car flow field features that can be observed, as the deceleration at the rear of the car caused by the presence of the rear wing. This phenomenon is also observed in the Minimum Subtraction, SFF and SFFM. Both Minimum Subtraction and SFF show regions of higher particle concentration at the front and rear of the car (see planes XZ and XY), caused by the reflections that appear on the surface of the car. The resulting velocity field contains areas that do not match their surroundings, leading to erroneous data in the velocity field.

In the case of SFFM, it shows similar results as the Minimum Subtraction and SFF, with a smaller FOV than the Butterworth's. Recall from the View 6 results, there were a few voids in the velocity field that appeared in all the planes showed, particularly in planes XZ (Figure 6.22) and YZ (Figure 6.22) with blank spaces caused by masking the rear wing's reflection. These gaps are no longer present in the merged-view results as they are successfully filled in by data from other views. However, there are still a few empty regions that are not properly addressed as occurs at the top of the rear wing in Figure 6.27 (fourth column), where there is a clear region that is successfully retrieved by the Minimum Subtraction, Butterworth and SFF, but not by the SFFM. Unfortunately, the measurement volumes acquired are not sufficient to fill in the gaps posed by the masking of this method. However, this could be fixed by acquiring additional measurement volumes with proper overlap between views.

The flow field at the top of the car and in the wake seem incomplete for the case of SFFM. As done for the side-view mirror, an intensity normalization is applied after the masking step to check if this increases the FOV. The results of this operation are shown in Figures 6.28 and 6.29.

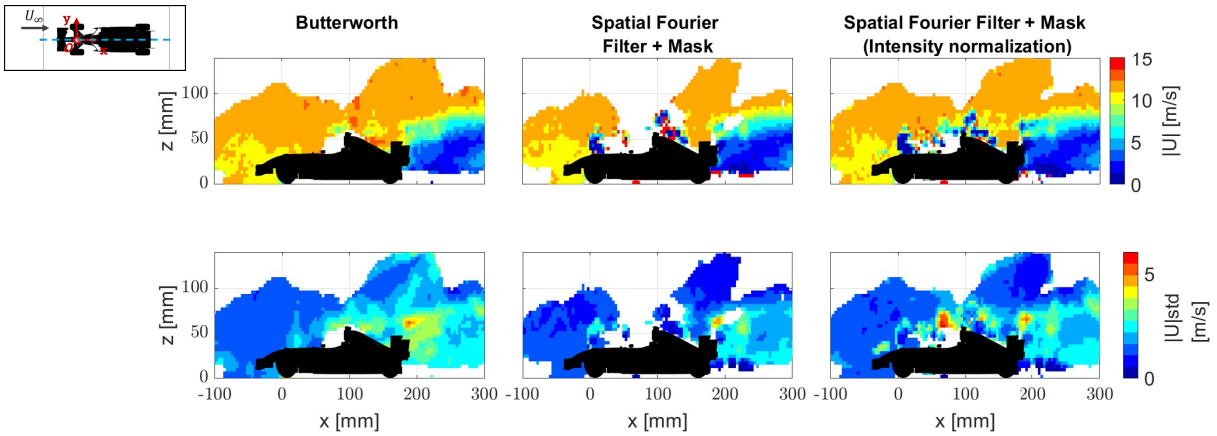

Figure 6.28: Formula 1's Butterworth, SFFM and SFFM with intensity normalization results on the XZ plane at y = 0 mm (all views).

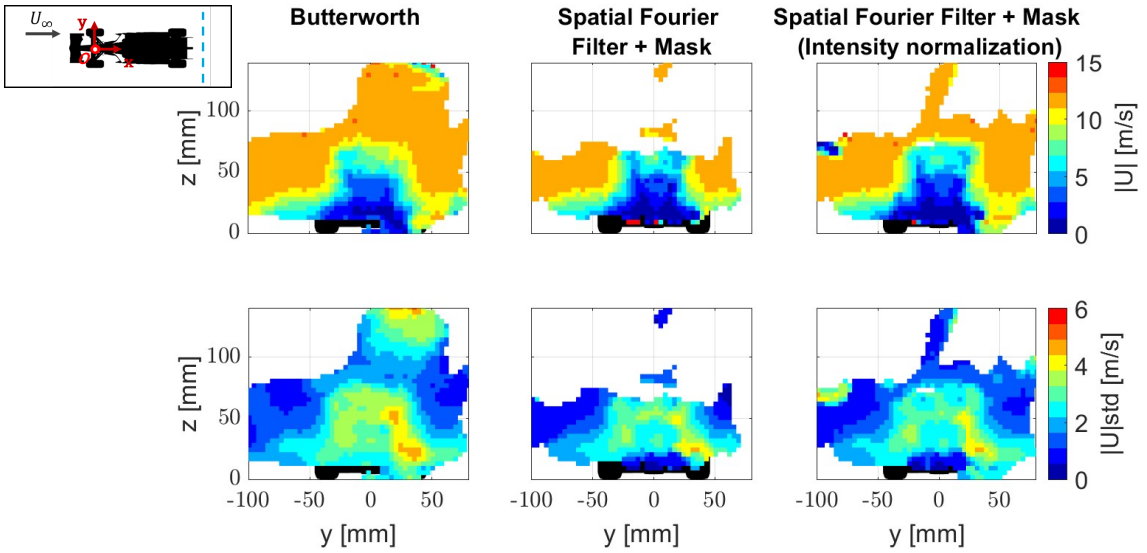

Figure 6.29: Formula 1's Butterworth, SFFM and SFFM with intensity normalization results on the YZ plane at x = 220 mm (all views).

The FOV is significantly increased, especially in those regions that the SFFM had voids in the data compared to the Butterworth. The top part of the car (near the exhaust) has been filled in with data, as well as the front and rear. However, there are still some non-removed parts of the reflection that can be identified in the flow field as clear outliers with higher standard deviation compared to the Butterworth's. All in all, adding an extra step of intensity normalization to the SFFM method does increase the amount of 3D data that is retrieved, but reflections regions that were not successfully removed have to be taken into account when examining the velocity field.

6.3. Propeller case

Finally, the last geometry tested is a propeller in pusher configuration and with a rotational speed of 3600 RPM. This geometry is expected to present unsteady reflections caused by the rotation of the blades and act as a crucial validation test case to check if the proposed methodologies are effective in mitigating unsteady spurious regions. The main focus of the robotic PIV measurements was the wake, trying to capture the propeller's streamtube features. The propeller in pusher configuration facilitates this, with the vertical and horizontal struts being upstream the region of the wake. Hence, the nine measurement volumes depicted in Figure 6.30 were acquired.

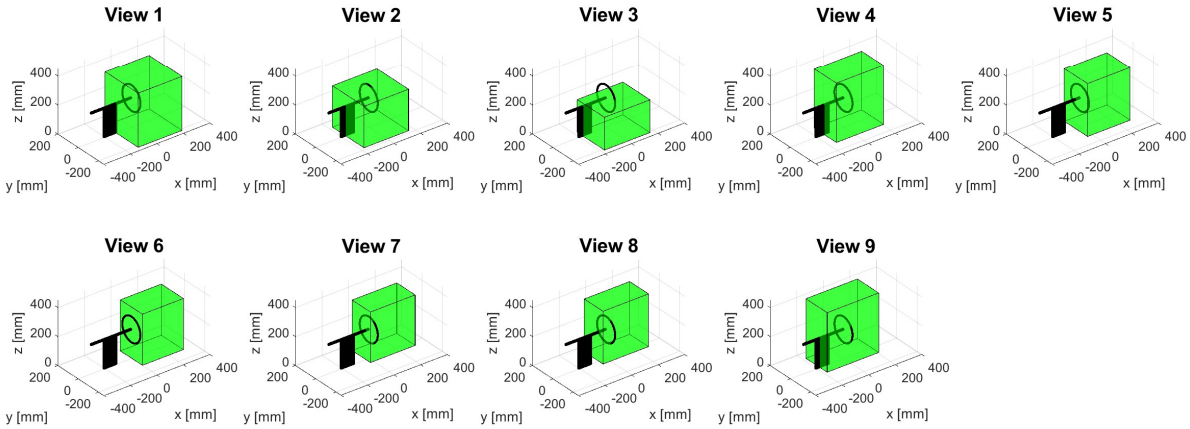

Figure 6.30: Set of measurement volumes acquired for the propeller.

It can be observed that all views capture the same region: downstream the propeller, on the wake. Most of them are centred on the propeller, except a few of them that point more towards the left-hand side (View 3), or the right (Views 7 and 8). What is really being recorded can be checked in Figure D.1 (see Appendix D) that shows an instantaneous of the raw and pre-processed images for each view. Not all views have dense reflections; for instance, in the case of views 6, 7 and 8 where there is no or little reflection in the images. In these cases, the Butterworth is able to remove almost all these regions. The problem arises with the rest of views, where the reflection on the propeller surface becomes more significant. There are three main sources of reflections: steady reflection from the vertical strut, unsteady reflection from the vibration of the horizontal strut and hub, and unsteady reflection from the blades. The one in the vertical strut can be easy to tackle by the Butterworth as it is steady. However, this technique fails to remove the other reflections that are fluctuating over time. One example of this is View 2, where the reflections are easily identified and occupy a big part of the image. The Butterworth case is able to eliminate the vertical strut's, but not the rest due to their fluctuating nature. The Spatial Fourier Filter is able to attenuate significantly all the regions with reflections, but, once more, there are still residual artifacts on these areas that do not get removed. When applying the extra step of masking in the Spatial Fourier Filter + Mask (SFFM), all the spurious regions are successfully and entirely removed from the images. This results in a clean-of-reflections set of images that will avoid the presence of spurious regions in the processed Shake-the-Box data.

A single instantaneous image is not sufficient to prove the ability of adaptive masking of SFFM. For this reason, a set of four consecutive recordings for views 2 and 4 are displayed in Figure 6.31 and 6.32, pre-processed with the Butterworth and Spatial Fourier Filter + Mask methodologies. The time filter fails to mitigate the intensity contribution corresponding to reflections on the moving objects. Conversely, this set of images proves the effectiveness of SFFM to generate a mask

specific for each recording. Notice that the shape of the mask changes for each recording, yielding a robustly elimination of all reflections that neither the Butterworth or the SFF could remove.

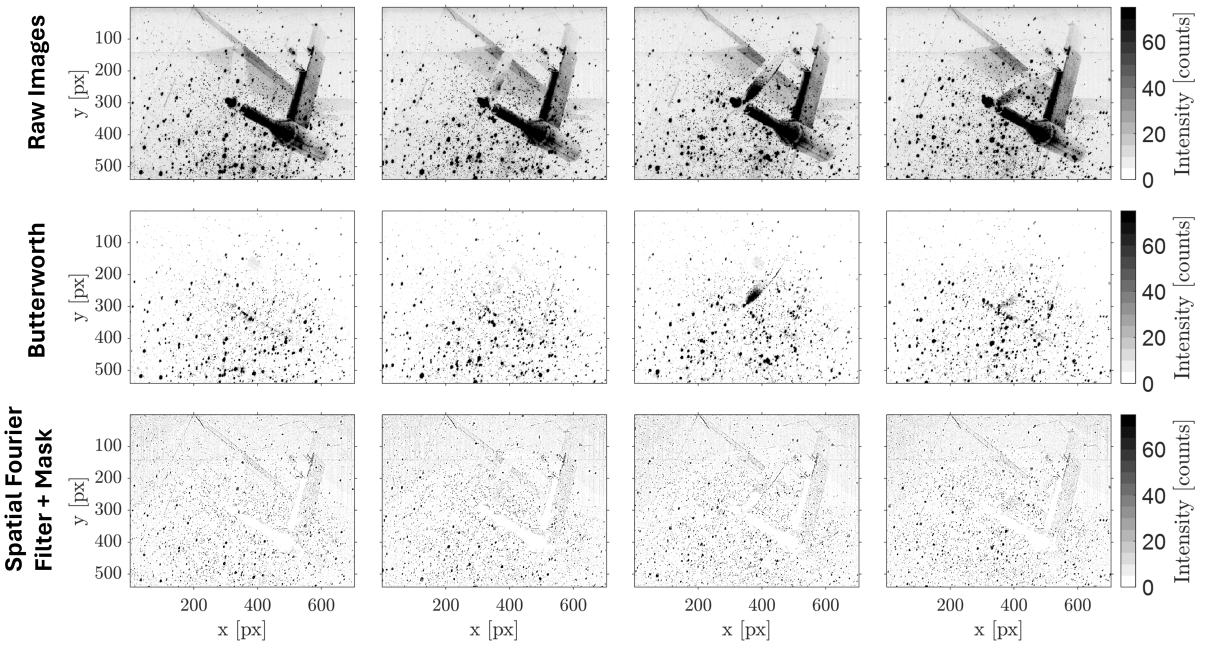

Figure 6.31: Set of consecutive pre-processed images (from left to right) with no filter, Butterworth time filter and Spatial Fourier Filter + Mask (View 2) of the propeller.

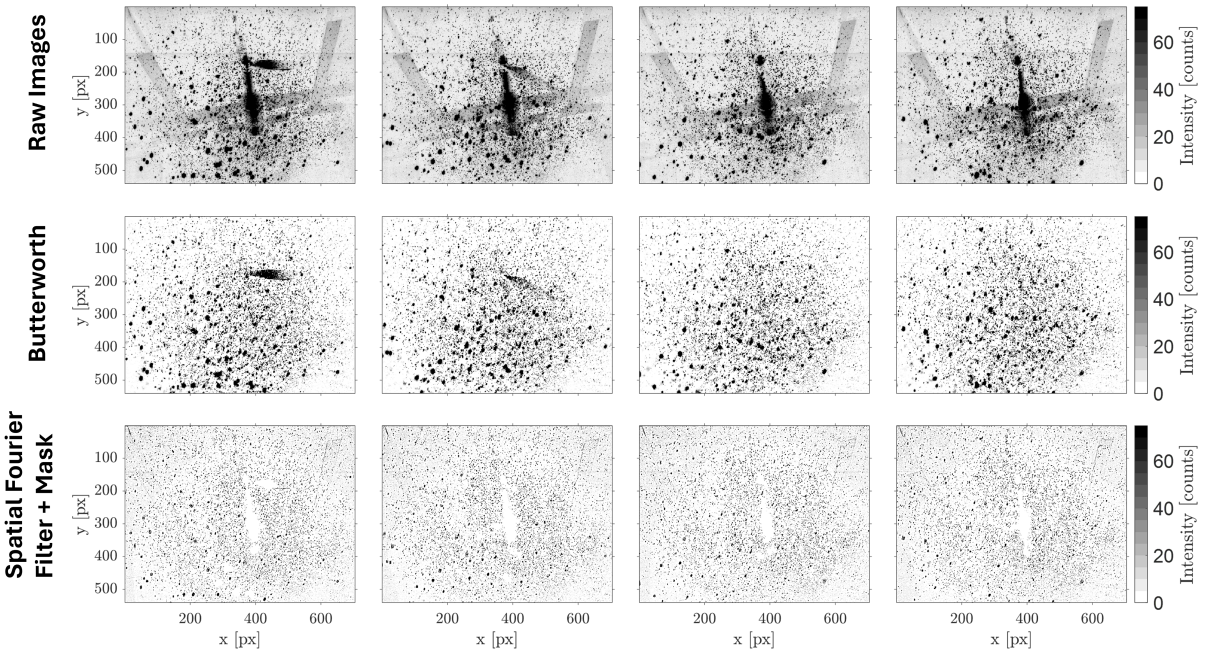

Figure 6.32: Set of consecutive pre-processed images (from left to right) with no filter, Butterworth time filter and Spatial Fourier Filter + Mask (View 4) of the propeller.

The number of particles tracked over time for each view and methodology is presented in Figure 6.33. First of all, the Minimum Subtraction fails no attenuate the reflections and, thus presents the highest number of particles detected (not all being real particles). Contrary to the trend observed for the side-view mirror (Figure 6.4) and Formula 1 car (Figure 6.20), here the Butterworth does not show the largest amount of number of particles and SFFM does not show the fewest N_p. The Butterworth, SFF and SFFM exhibit similar values of particles detected for several views (5, 6, 7 and 8), where the resulting images processed with the three approaches are similar in terms of reflection being attenuated and the amount of particles kept (see Figure D.1). In the rest of the views, there is a noticeable difference of N_p, with SFFM having the lowest value, result of applying a mask to the images. Finally, regarding the Particle Concentration Mask, from view 4 to 9 it is able to identify regions of reflection and remove them since it shows less number of particles detected than the Minimum Subtraction (Particle Concentration Mask masks the result from this approach).

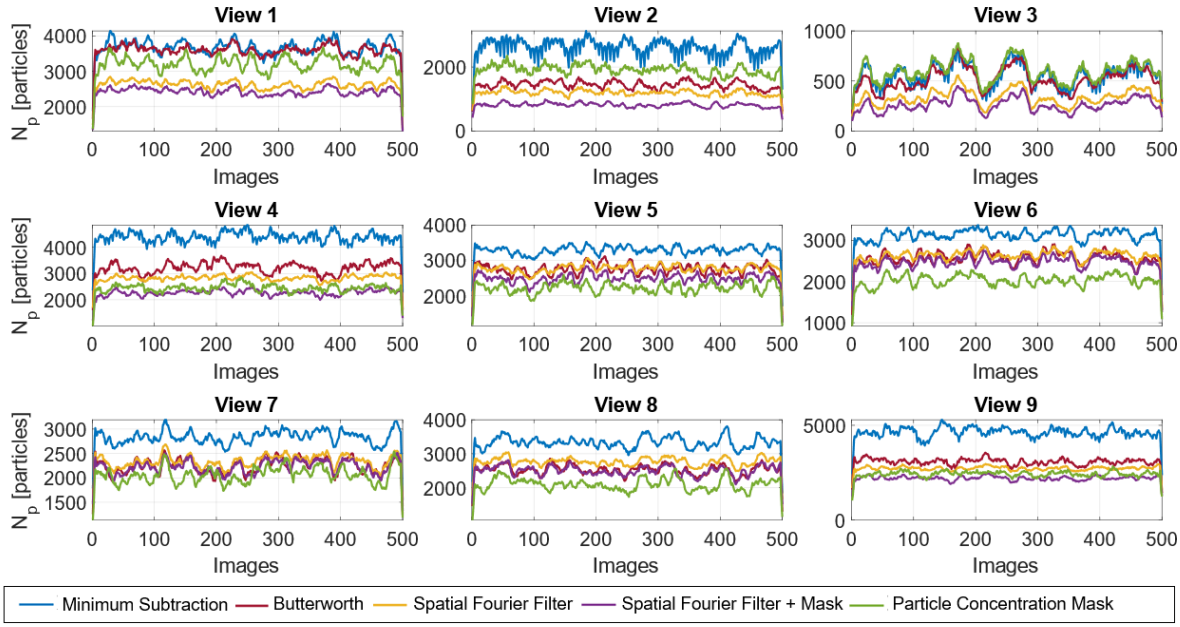

Figure 6.33: Number of particles per recording for each view (propeller case).

6.3.1. Individual views

In order to see how each methodology performs in a single view, the measurements volumes 2 and 4 are examined in this section (examples of views with the largest reflections). Figure 6.34 show the raw and pre-processed images of View 2.

Figure 6.35 shows the YZ plane at the wake of the propeller (5 cm downstream) of View 2. In all cases, the streamtube from the propeller can be identified, presenting a velocity increase from 5 m/s to approximately 8 m/s. The Minimum Subtraction and Particle Concentration Mask do not remove the reflection caused by the blade, producing the appearance of an artifact with the shape of the blade in the velocity field. This phenomenon also occurs in the Butterworth, although smaller in size. This does not happen in the case of SFF and SFFM, with the latter showing a blank due to the mask. The reflection caused by the propeller struts can be observed on the XZ and XY planes in Figures 6.36 and 6.37, respectively. All cases, except the SFFM, show a region of higher particle concentration, mean velocity and standard deviation that is due to the struts' reflections. This does not appear in the case of SFFM, where there is a void in the data, which is expected to be filled in by other views. Additionally, the latter approach exhibits the lowest standard deviation in all planes (indicator of good and reliable results).

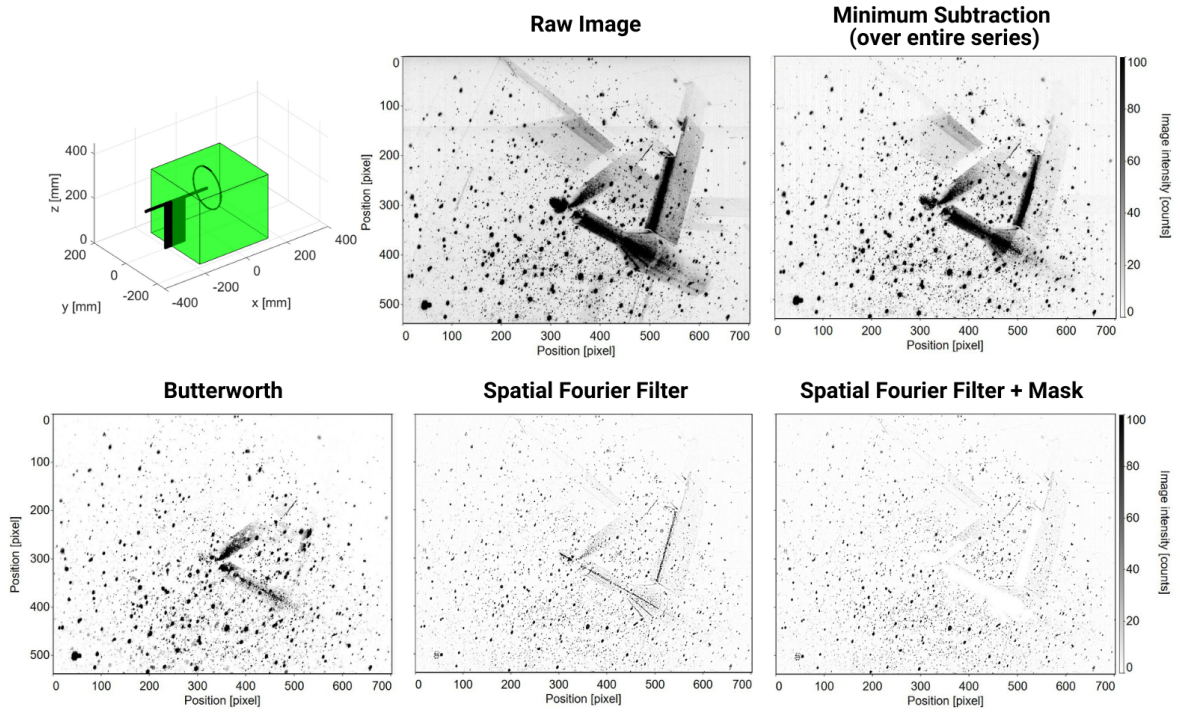

Figure 6.34: Propeller acquired and pre-processed images of View 2.

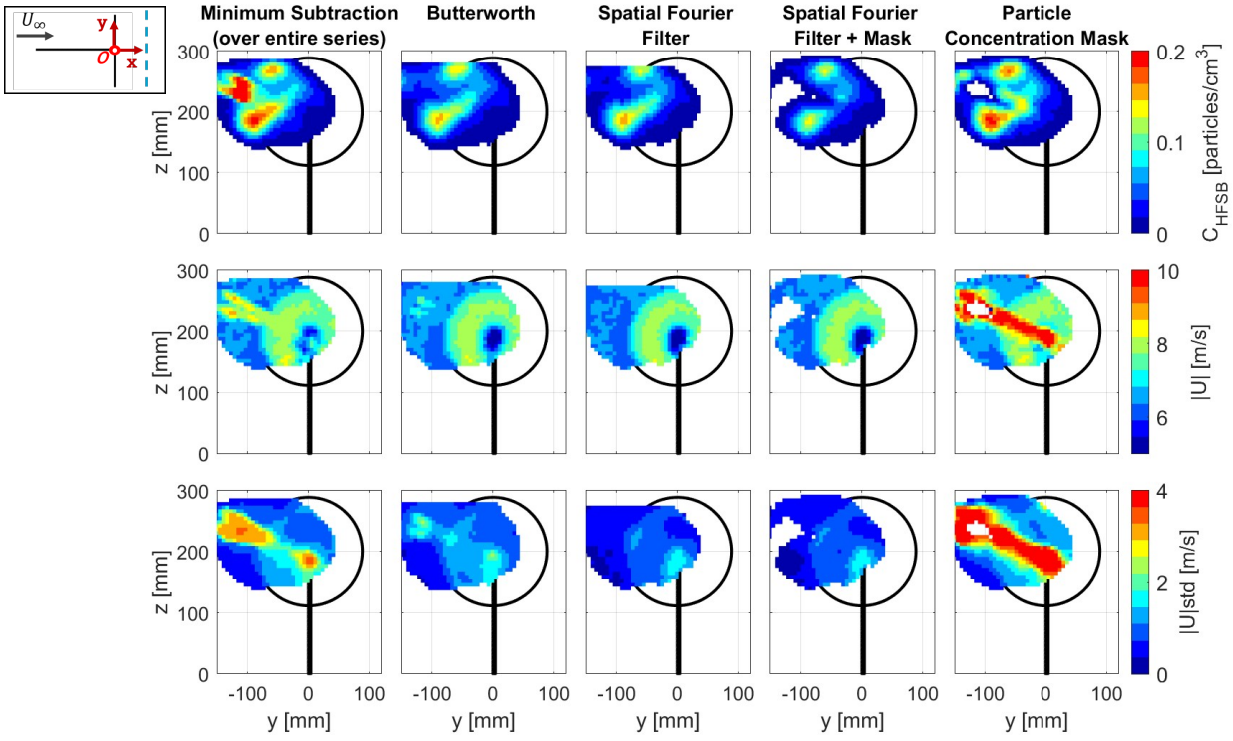

Figure 6.35: Propeller binning results on the YZ plane at x = 50 mm (View 2).

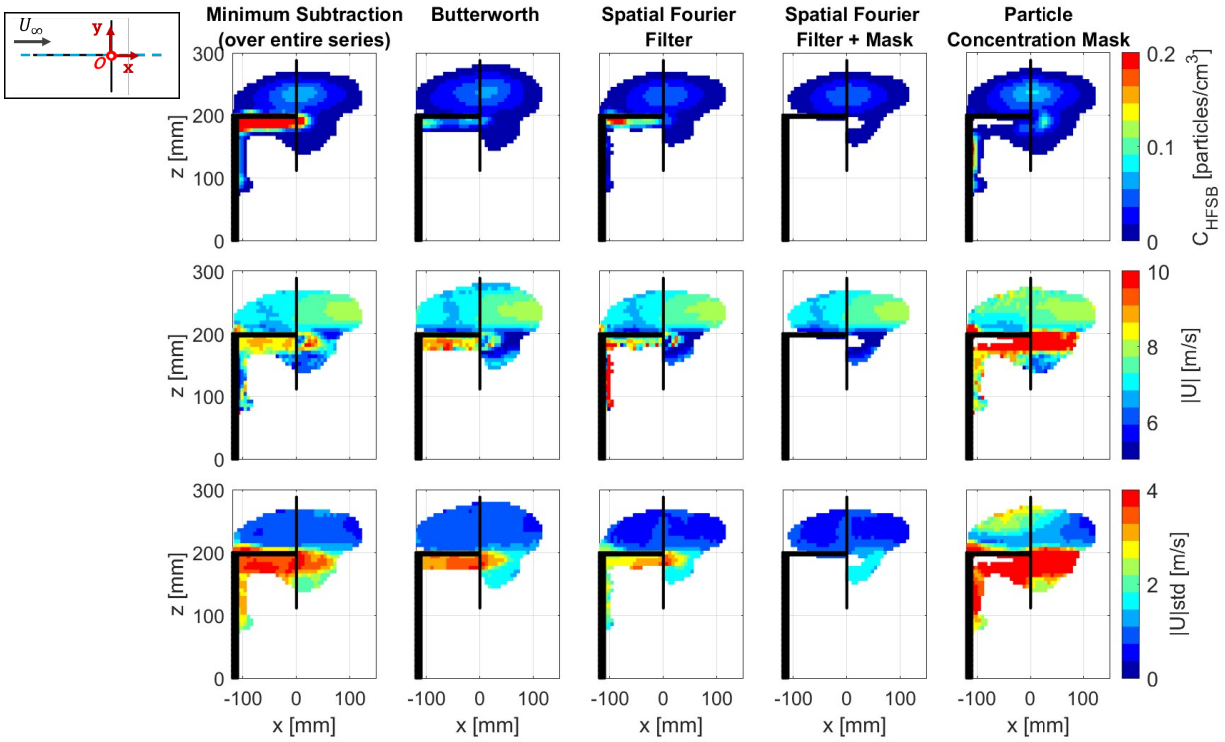

Figure 6.36: Side-view mirror binning results on the XZ plane at y = 0 mm (View 2).

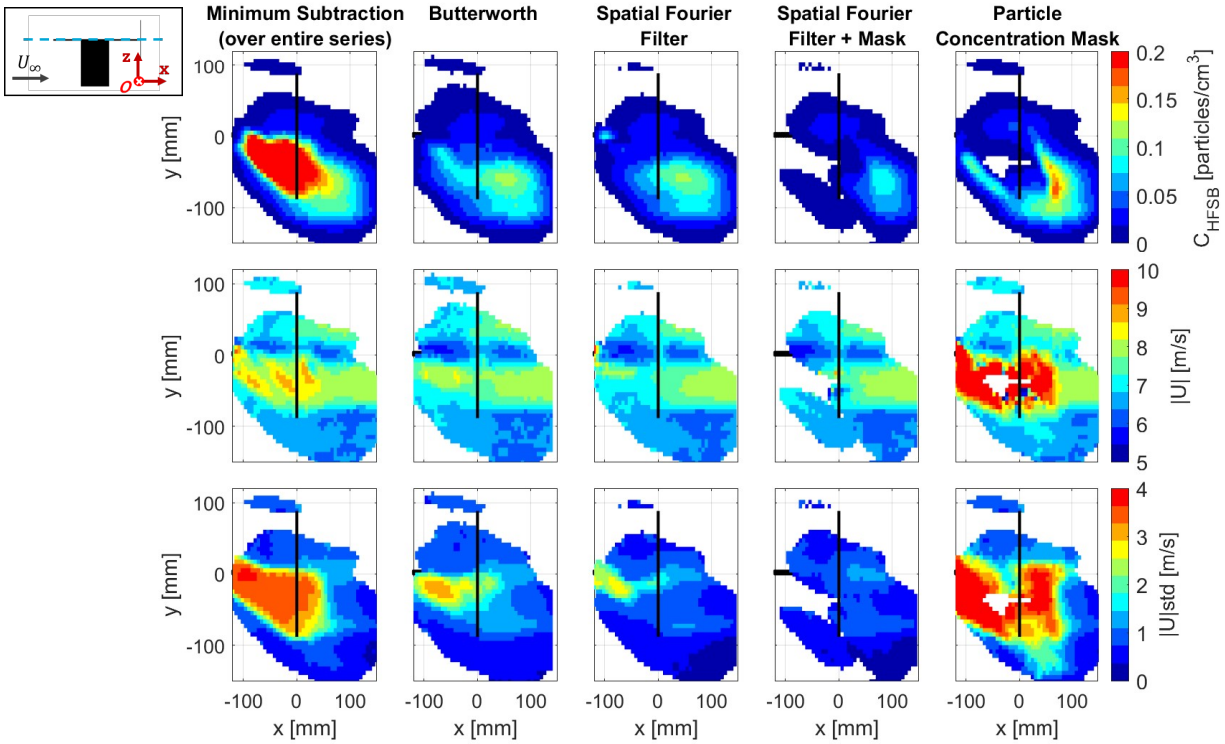

Figure 6.37: Side-view mirror binning results on the XY plane at z = 200 mm (View 2).

The next measurement volume to examine is View 4 with its corresponding images shown in Figure 6.38. This view presents a steady large reflection from the propeller strut and an unsteady one originated by the blade's rotation. As discussed previously, the Minimum Subtraction result does not deviate much from the raw image, while the Butterworth and SFF attenuate the reflection significantly, but not completely. Then, thanks to the additional masking step, the SFFM successfully removes the existing reflection resulting in a void in this region, where also no particles will be tracked.

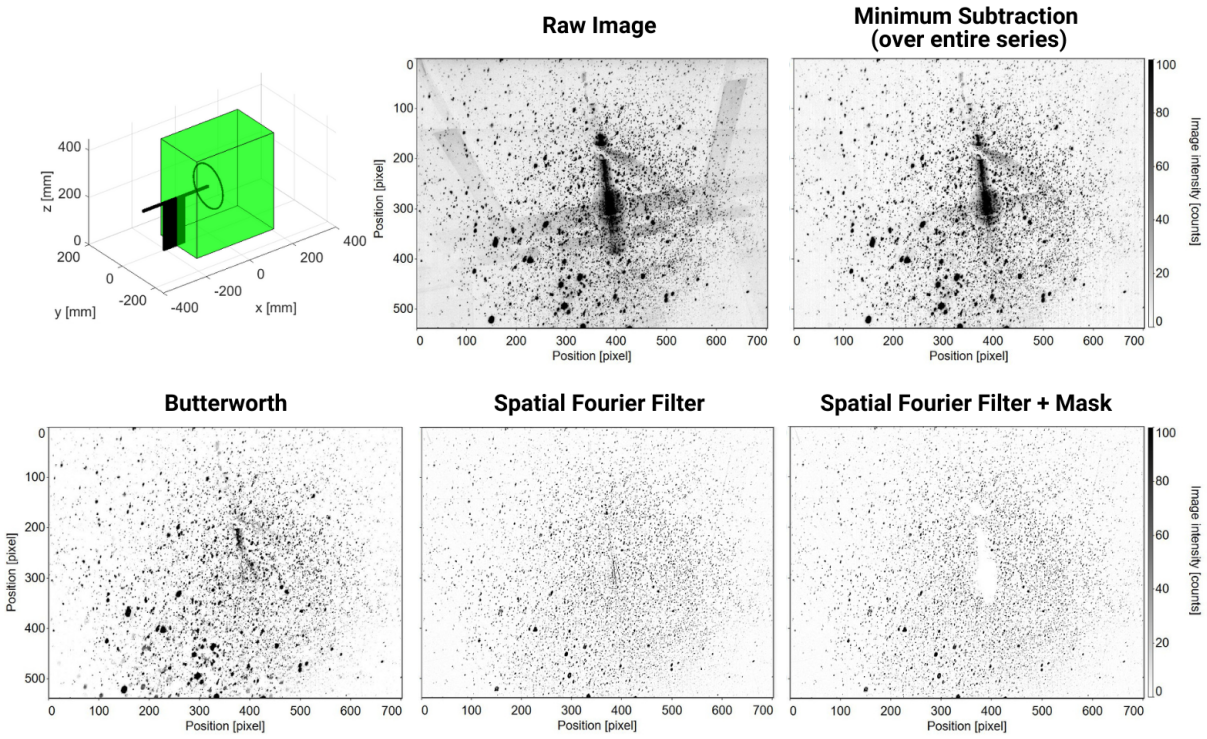

Figure 6.38: Propeller acquired and pre-processed images of View 4.

Figures 6.40 and 6.39 show the results at XZ and YZ planes of View 4. Similarly to View 2, the reflection caused by the blade and the strut appear in the Minimum Subtraction, Butterworth and Particle Concentration Mask resulting flow field. For the latter approach, it can be observed how it manages to mask the areas with higher particle concentration, but it still does not translate into a correct velocity field. SFF and SFFM show the propeller's streamtube with the corresponding velocity increase and a low standard deviation. Moreoever, in the case of SFFM there is a gap in the data caused by the propeller's strut mask.

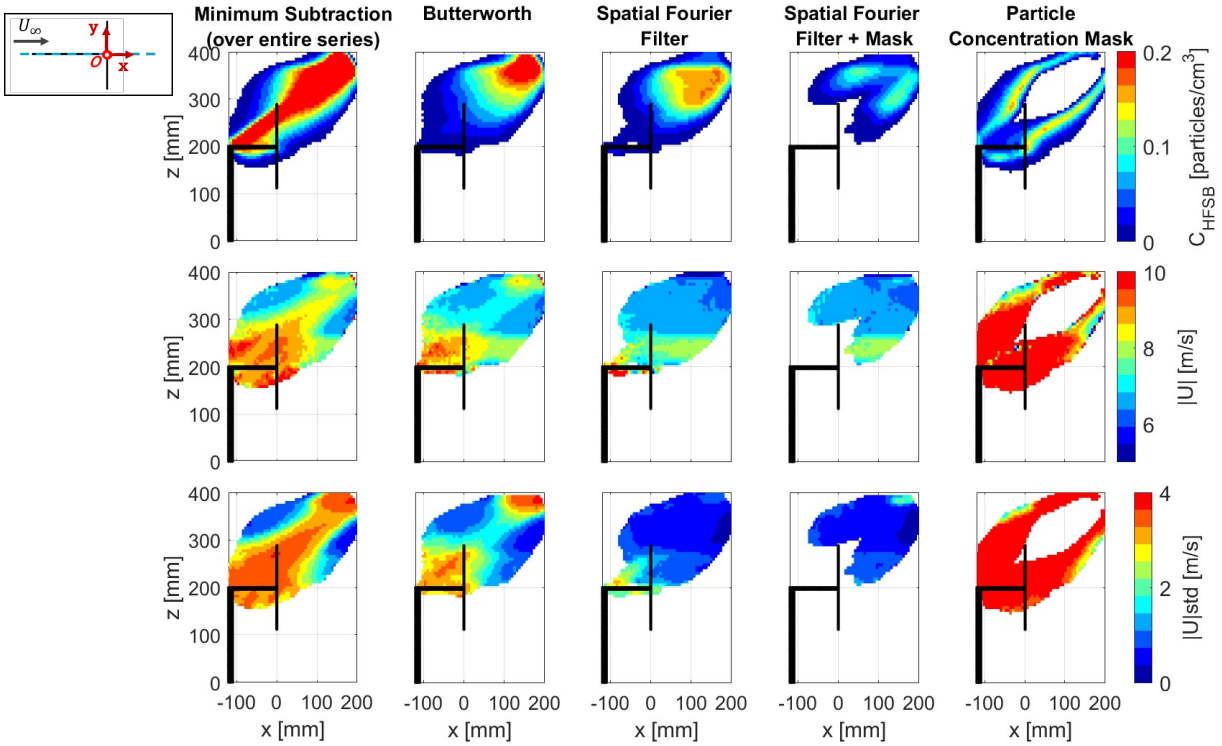

Figure 6.39: Side-view mirror binning results on the XZ plane at y = 0 mm (View 4).

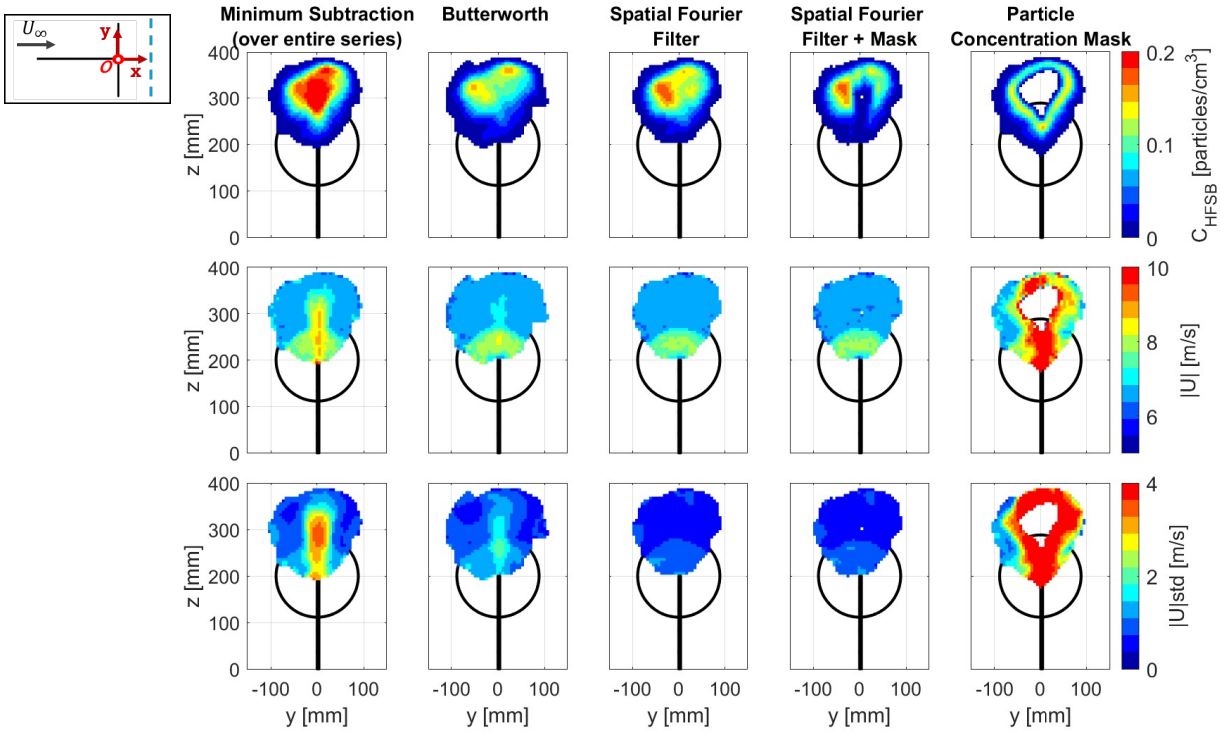

Figure 6.40: Propeller binning results on the YZ plane at x = 50 mm (View 4).

6.3.2. Complete measurement domain

Once processed all measurement volumes separately, the complete resulting flow field is obtained by performing the merging operation of the whole set of views. Then, the binning operation is executed using the same parameters as those applied to the individual views. The particle concentration, mean velocity and mean velocity standard deviation of the merged dataset can be checked in Figures 6.41, 6.42 and 6.43 for the YZ (wake), XZ (propeller's symmetry plane) and XY planes, respectively.

The YZ plane (Figure 6.41) shows the flow field 5 cm downstream, at the wake of the propeller. Almost the entire streamtube can be observed in the case of the Butterworth, SFF and SFFM. The part missing may be due to that region being less accessible by the robot. The Butterworth result has a slightly larger field of view than the SFF and SFFM, but the standard deviation is lower in the latter two cases. In the XZ plane (Figure 6.42), a side view of the flow field at the propeller's symmetry plane can be analysed, and the XY plane (Figure 6.43) shows a top view of the propeller with respect the propeller's hub (where the blades are attached). Both these planes clearly show the wake from the propeller's streamtube. The Minimum Subtraction, Butterworth and Particle Concentration Mask's flow fields present the influence of the blade and strut's reflection that propagate along the line-of-sight when processing with the Shake-the-Box algorithm. SFF shows similar results to SFFM, but has higher standard deviation, which indicates worse quality of the results. Additionally, it is important to notice that the voids that appeared in the individual views are successfully filled in in the merged-view case.

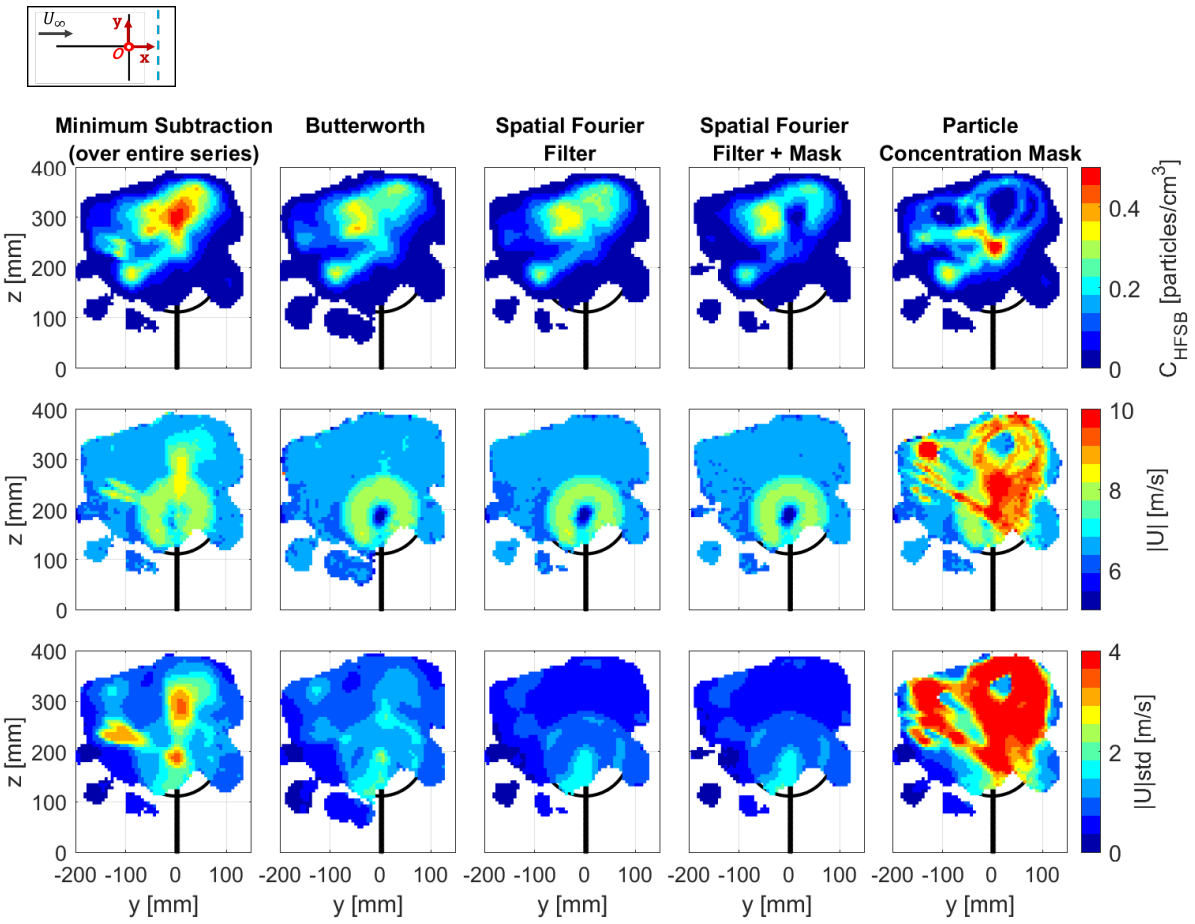

Figure 6.41: Propeller binning results on the YZ plane at x = 50 mm (all views).

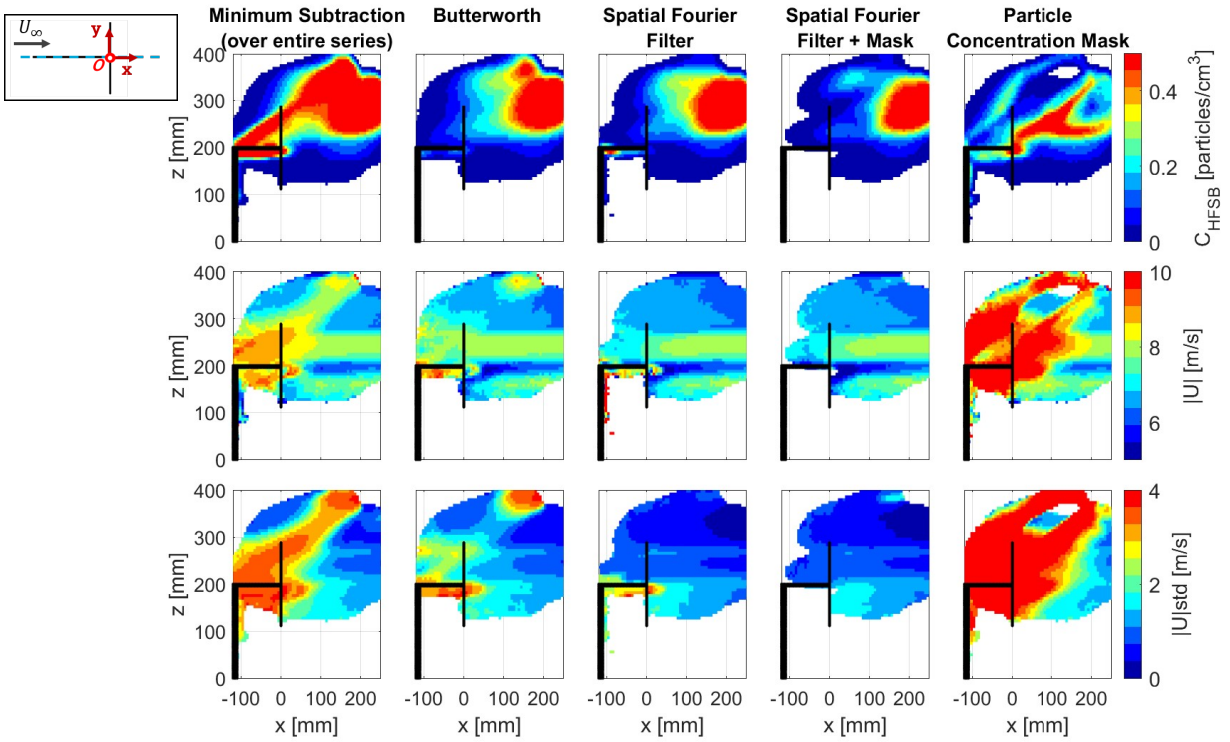

Figure 6.42: Side-view mirror binning results on the XZ plane at y = 0 mm (all views).

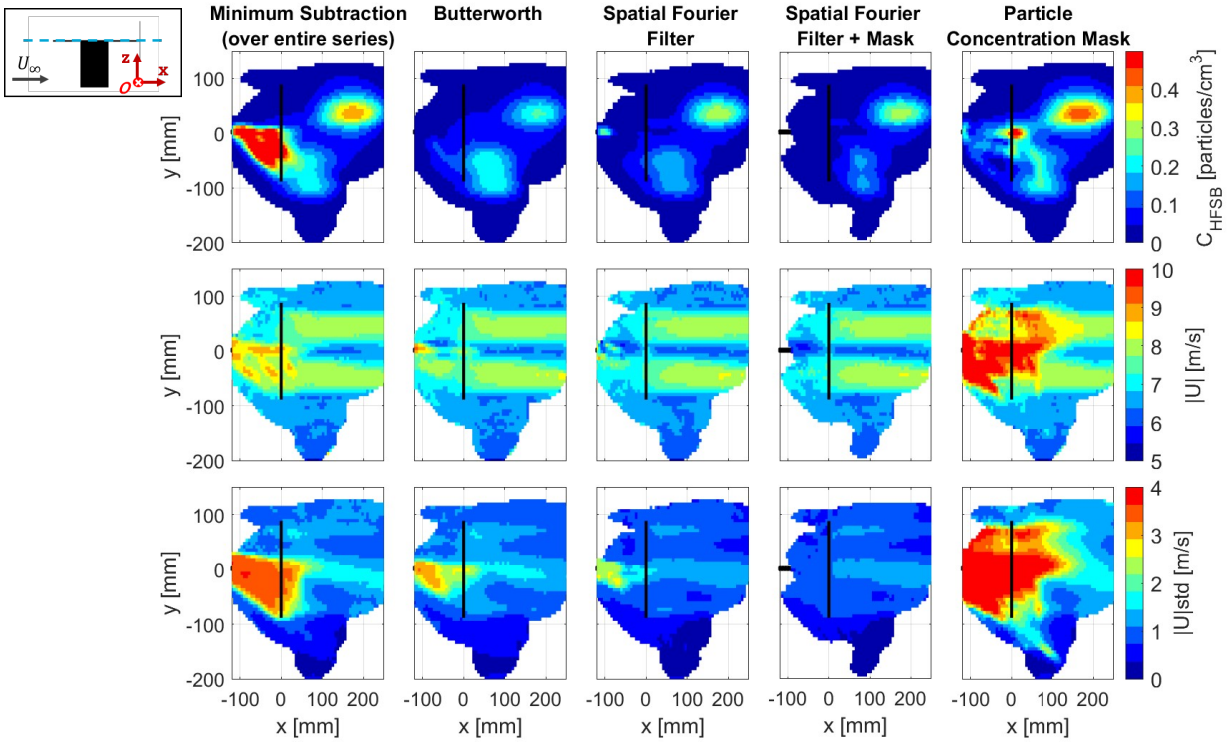

Figure 6.43: Side-view mirror binning results on the XY plane at z = 200 mm (all views).

As occurred with the side-view mirror and Formula 1 car, it is important to notice that the field of view of the SFFM is slightly smaller than the Butterworth. Even though, the time filter was not effective in removing the reflections, the particles that remained in the pre-processed images are real and they are the reason why this approach results in a larger FOV. For this reason, the SFFM results of the propeller are also tested with the implementation of an intensity normalization with a local average over 100 pixel.

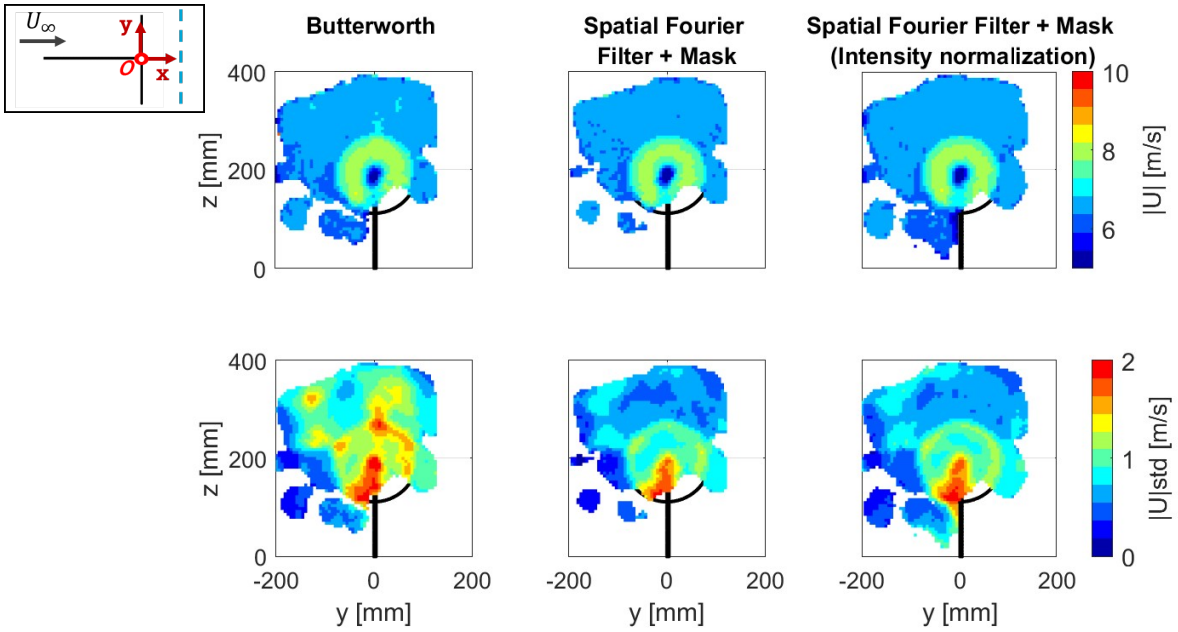

Figure 6.44: Propeller Butterworth, SFFM and SFFM with intensity normalization results on the YZ plane at x = 50 mm (all views).

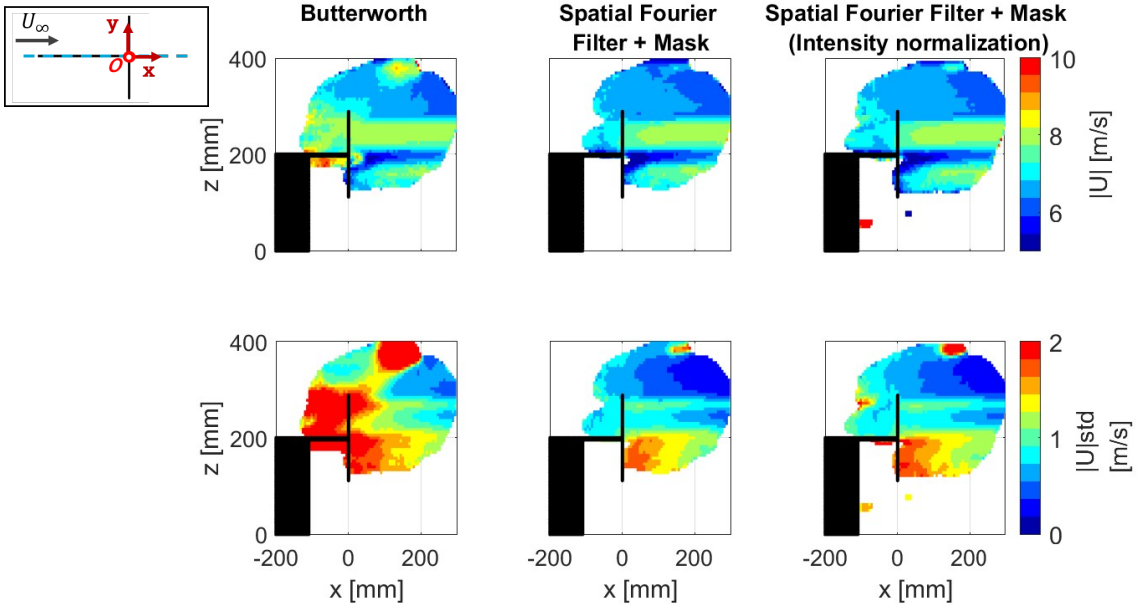

Figure 6.45: Propeller Butterworth, SFFM and SFFM with intensity normalization results on the XZ plane at y = 0 mm (all views).

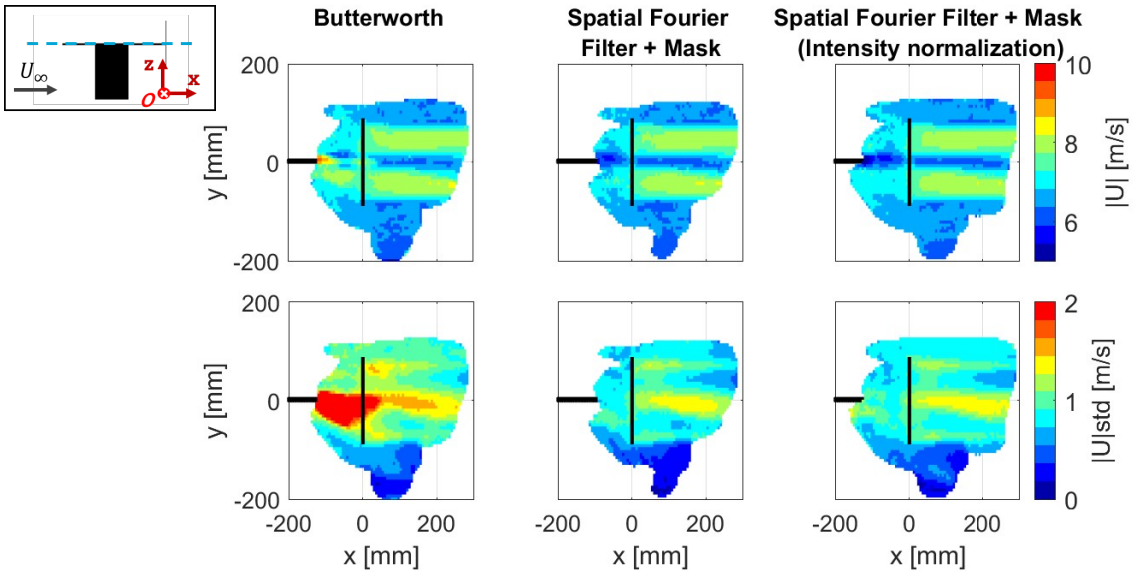

Figure 6.46: Propeller Butterworth, SFFM and SFFM with intensity normalization results on the XY plane at z = 200 mm (all views).

An important implication is that the reduction in FOV of SFFM compared to the Butterworth filter is primarily due to a non-uniform intensity distribution across the image. This issue is easily addressed by performing intensity normalization on the entire set of images as proven in Figures 6.44, 6.45 and 6.46, thus allowing for the recovery of the lost FOV, while maintaining low standard deviation values (indicator of good quality of the results).

6.3.3. Further discussion on Butterworth - SFFM

The Butterworth time filter and the proposed masking approach are the two methodologies that exhibit the best results. Therefore, this section presents the results of processing the whole set of 5,000 images with STB with these two strategies. Figures 6.47 and 6.48 show the binned velocity field on YZ and XY planes, respectively. The results are given for three single views as well as the merged dataset. For the latter, the standard deviation of the velocity magnitude is also given in the figure.

The most noticeable feature is that the spatial gaps that appeared in the previous results of only processing 500 images are completely filled. This is thanks to the amount of data present in the set of 5,000 images. Hence, it is important to consider the number of images to acquire when implementing the spatial Fourier filter and mask approach.

Figure 6.47 shows the resulting velocity field on the YZ plane for two single views and the merged-view case, showing the flow field 5 cm downstream, at the wake of the propeller. The Butterworth attenuates the reflection, but not completely as the rotating blade shape appears in the domain due to the residual artifacts not removed on this region (as shown in Figure 12), which propagate along the line-of-sight of the velocimeter when processing with the Shake-the-Box algorithm. The spatial Fourier filter and mask velocity field presents regions that are not captured in the single views compared to the Butterworth's: bottom streamtube in View 1 and bottom and top right in View 2. However, as expected, these gaps are successfully filled in when merging all the views together. The entire streamtube can be observed in both cases, exhibiting a velocity increase from 7 m/s to 8 m/s. The additional masking step allows to successfully remove the reflection from the blades, showing a lower mean velocity and standard deviation. The latter is the most distinguishable feature of this method, showing a much cleaner and reliable flow field.

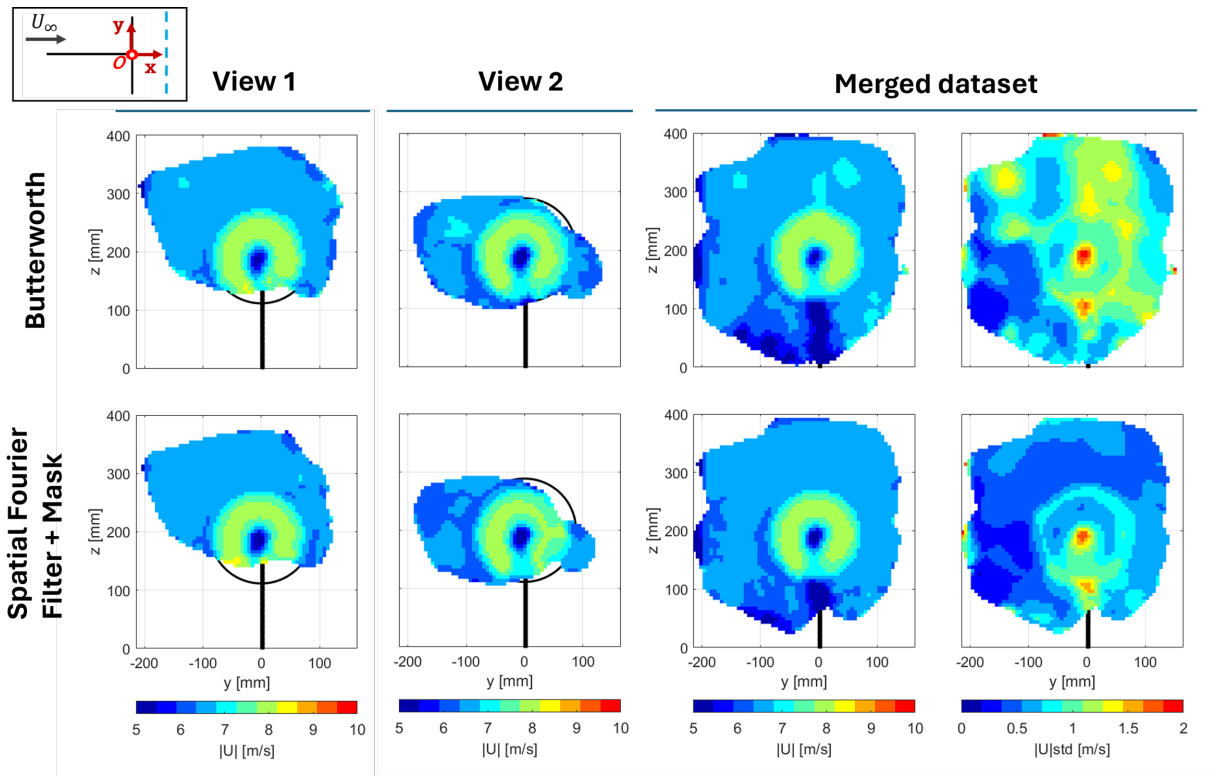

Figure 6.47: Propeller binned results on YZ plane at x = 50 mm for single and merged views (Butterworth and Spatial Fourier Filter and Mask).

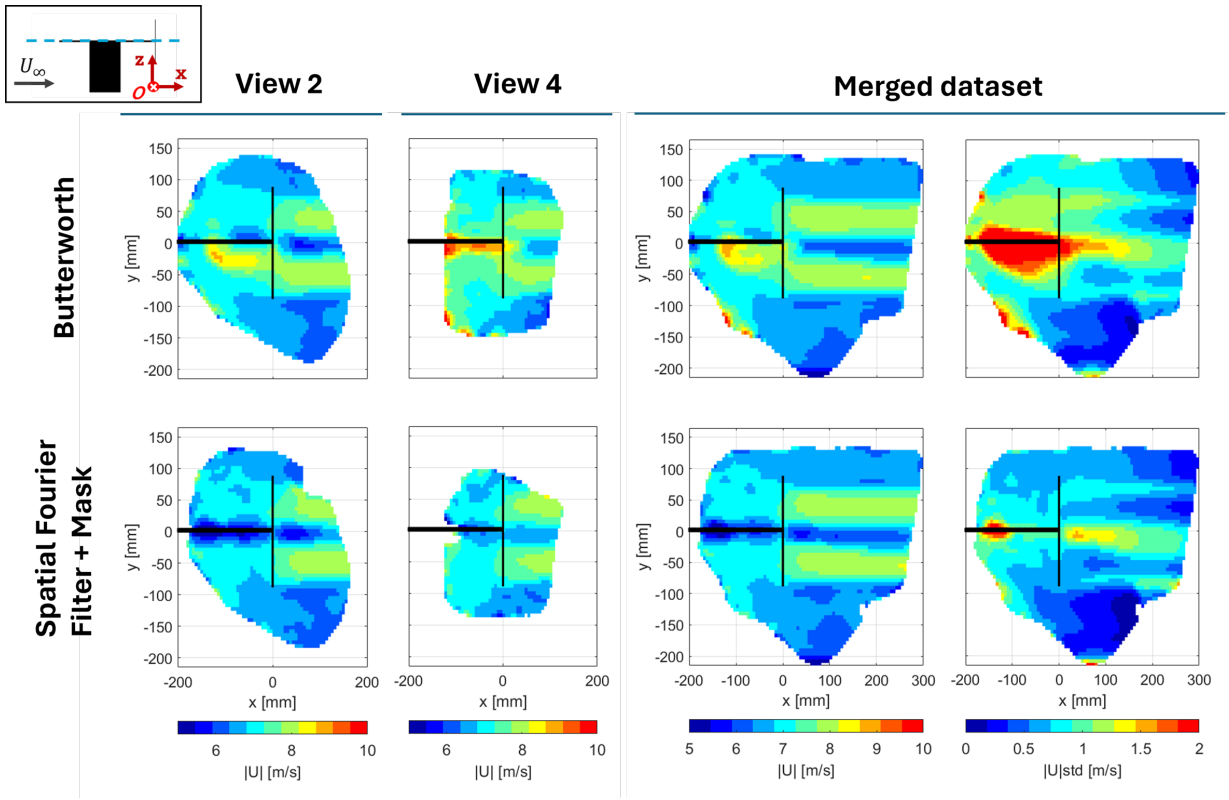

Figure 6.48: Propeller binned results on XY plane at z = 200 mm for single and merged views (Butterworth and Spatial Fourier Filter and Mask).

The Butterworth flow field shows the influence of the propeller horizontal strut reflection, which appears as a high mean velocity region as observed in View 4 (Figure 6.48). This does not occur for the masking method, which shows a small gap, consequence of masking this area. Since View 2 contains data in the region within this void, the latter is entirely filled in the merged-view case. For the complete velocity field, the mean velocity and its standard deviation is shown. It is important to highlight that the spatial Fourier filtering and masking approach has a significantly lower standard deviation compared to the Butterworth, indicating much more reliable results. This is thanks to the successful removal of both steady and unsteady reflections with the masking approach. Conversely, the Butterworth shows a significant high standard deviation in the region of the horizontal strut as a result of the reflection that emanates from this part of the geometry.

7

Conclusions and recommendations

In this chapter, a summary of the study's key findings is provided along with the discussion of their implications, reflecting on the significance of the research and its potential impact on the field. It also identifies areas for further research and suggests possible directions for future work based on the insights gained from the analysis.

7.1. Conclusions

This study introduced and evaluated three novel approaches for mitigating laser light reflections in instantaneous data from Robotic Particle Image Velocimetry (PIV) experiments. These methods aimed to address the challenge of high-intensity reflection regions corrupting particle detection and analysis in volumetric PIV systems, particularly in Robotic PIV.

The first and second methods involved image filtering and masking techniques in the wavenumber space, by using the 2D discrete Fourier transform (DFT) to decompose the image signal into low- and high-wavenumber components. In the first method named *Spatial Fourier Filter*, a high-pass filter is applied to attenuate the intensity of reflection regions composed mainly by low-wavenumber content. Then the second method called *Spatial Fourier Filter + Mask* was introduced, following the same process as the first method with an additional automated adaptive masking step to remove residual reflection areas that the filtering could not eliminate.

In the case that none of the existing or proposed image pre-processing methodologies work, the third methodology *3D-based Particle Concentration Mask* applied on a later stage of the processing pipeline is proposed. This creates a 3D mask on the instantaneous processed Shake-the-Box data by analyzing the particle concentration distribution over the flow domain and applying a threshold with a selected reference value.

These approaches were tested and compared to state-of-the-art techniques on experimental data obtained from Robotic PIV experiments conducted on three different geometries: a side-view mirror, a Formula 1 car, and a propeller. The results demonstrated the effectiveness of the proposed approaches in eliminating reflection regions and preventing the appearance of spurious particle tracks. Comparison between raw and pre-processed images, as well as particle tracking results, confirmed the successful removal of reflection-induced artifacts using the spatial Fourier filter automated masking approach. The outcome of the experiments is addressed by revisiting the research questions.

The research question and subquestions posed in Chapter 1 are repeated below for convenience; RQ refers to *Research Question* and RSQ refers to *Research Subquestion*.

(RQ.) How can a methodology effectively identify and mitigate the effects of spurious regions of reflections in Robotic PIV across various PIV datasets?

(RSQ1) Which image pre-processing technique can be applied to detect and mitigate reflections by only using an instantaneous image?

The literature review revealed several state-of-the-art techniques for reflection mitigation. However, most of them were mostly effective in cases of steady reflections, like the POD [65] and Butterworth time filter [64] approaches. There is one methodology proposed by Adatrao et al. [66] that deals with unsteady reflections, but just when these are sharp. These would not be effective removing reflections from Robotic PIV images as they generally appear as glare spots. By inspecting the reflection characteristics compared to those of the particle tracks, one can see that there is a potential way to distinguish between both objects by filtering or classifying the image objects by their size or wavelength. Reflection larger in size by nature have longer wavelengths, while particles being smaller (around a few pixels) present shorter wavelengths.

Therefore, here there is the possibility of filtering out by wavelength size or in other words, by wavenumber (inverse of the wavelength) content by using spatial Fourier analysis. This follows a similar methodology of the Butterworth time filter, but instead of filtering in the time frequency domain, it is done in the spatial frequency or wavenumber domain. Spatial Fourier filtering is a common technique used in image pre-processing that can have several applications: using a low-pass filter to smooth the images or highlight the edges of an image by applying a high pass filter. For processing the PIV acquired images and only keeping the particles, a high-wavenumber filter must be used. This leaves out the low wavenumber content corresponding to big artifacts in the images, like reflections, and keeps the high wavenumber objects, like particles. This proposed methodology is called *Spatial Fourier Filter* (also referred to as SFF).

As a further step, one can take advantage of the reflection attenuation to identify the regions in the image that have suffered the most intensity variation after applying the filter. Then, after identifying these areas, a mask around them can be created with the aim to remove all the components of the reflection for every single instantaneous recording, leading to an automated adaptive masking approach denominated as *Spatial Fourier Filter + Mask* (or SFFM).

(RSQ2) How do reflections appear in the 3D velocity field? Is there a distinguishable characteristic that could be attributed to and, hence used to remove this region of the data?

A reflection characterization has been done in both PIV images and the processed Shake-the-Box results. Reflections appear in the 3D velocity field as regions with distorted or erroneous velocity information. They often manifest as accumulations of particles that propagate along the line-of-sight of the camera. These particles, known as ghost particles or false positives, appear within the reflection region due to their similar intensity to the particles. As a result, the particle tracking algorithm misinterprets these pixels as particles, even though no particles are actually present.

When examining the velocity information within reflection regions, it is challenging to distinguish between real particle tracks and reflections. Particles detected inside reflection regions can exhibit a wide range of velocities, varying from high to low. Therefore, analyzing velocity information alone does not provide a clear distinction between particles and reflections. However, when

analysing the particle concentration over the flow field domain, the regions of reflection show higher particle density values compared to the rest of the flow. The Robotic PIV data acquired for a sphere geometry by Jux et al. [86] is used as example to test if the methodology hypothesis is correct. In this test case, there is a clear reflection that emanates from the surface of the sphere, which then appears as an accumulation of particles in the processed 3D velocity data. As expected, this region also appears as a higher particle concentration area compared to the rest of the flow field. Hence, by selecting a particle concentration reference value, the dataset can be masked such that regions with higher or equal concentration than the reference value are removed, and regions with lower values are kept.

The following two subquestions are addressed simultaneously, with each of the proposed methodologies discussed separately for the reader's convenience.

(RSQ3) What are the advantages of the developed methodologies against state-of-the-art techniques in identifying and mitigating spurious regions?
(RSQ4) What are the potential challenges and limitations in implementing the proposed methodology and how can these be addressed?

To study the implications and differences of the three proposed reflection treatment methodologies, the later are compared with two state-of-the-art techniques: the Minimum Subtraction over the entire series and the Butterworth high-pass time filter.

First, recall the reflection characteristics on each of the geometries tested. The side-view mirror presented several reflection across the set of views acquired. The most significant came from the floor and the foam attached on the surface of the mirror. In the case of the Formula 1 car, the reflections appeared in a different manner: they adopt the shape of the car and its little details (e.g. the front wing or front suspensions). And finally, the propeller posed the most challenging case with reflections originating form three different sources: steady ones from the propeller's vertical and horizontal struts, and unsteady ones caused by the blades' rotation. As the last one is moving over the entire set of images, its consequential reflection does not appear in all images, but indeed depends on the rotation of the blade and its position with respect the CVV probe.

Spatial Fourier Filter (SFF)

The Spatial Fourier Filter is able to attenuate the regions of reflections from each instantaneous recording in all test objects, reducing the size of unwanted regions. However, it has an important drawback: it fails to remove the high-wavenumber components present in reflections. This is because reflections are not perfectly uniform, but they show some granularity, which makes the reflection not only contain low-wavenumber, but also high-wavenumber content. Consequently, residual reflection components are still visible in the SFF images. This is particularly noticeable in the case of the Formula 1 car, where the resulting velocity field shows nonphysical regions. Despite this limitation, this approach shows a lower standard deviation compared to the state-of-the-art methods, which is an indicator of good flow field quality.

Spatial Fourier Filter + Mask (SFFM)

The clear limitation posed by SFF can be overcome by the Spatial Fourier Filter + Mask approach, where the regions of reflections are identified and masked out, yielding an resulting image and 3D processed velocity field free of reflections. This method does the same as SFF, but with and additional step of automated adaptive masking; the areas that suffer the largest variation with respect the original image are considered potential artifacts to be masked. This is clearly proven in the resulting images of the side-view mirror, Formula 1 car and propeller. The steady reflections from the first two cases are properly masked, yielding a velocity field similar to the Butterworth's. Note that in this case of steady reflections, the Butterworth results are considered as ground truth.

Hence, both Butterworth and SFFM are successful at removing all spurious regions, but in the case of SFFM, it additionally shows a lower standard deviation within the flow field.

The propeller case supposed a challenging case, crucial to determine the effectiveness of SFFM in reflection mitigation of unsteady reflections in individual instantaneous images. As stated before, most of the propeller's views showed a big reflection coming from the strut and a moving reflection from the blade. Also, the horizontal strut suffered a slight vibration due to the blades' rotation, which supposed an additional challenge for the state-of-the-art techniques. The Minimum Subtraction failed at removing any of the artifacts, and the Butterworth time filter could remove part of the steady reflections from the vertical strut, but could not manage to eliminate the others. Here, the SFFM makes a difference, being able to create a mask of all existing reflections and mitigate their contribution to the 3D Shake-the-Box results. Furthermore, a set of consecutive images for different views have been displayed proving the method's capability to generate an adaptive mask for each of the instantaneous images, adopting the shape of the spurious regions and effectively capturing and removing the reflection on the rotating blade.

The potential side effect of masking an image is that it leaves a empty space without reflections, but also without any particle tracks that could have crossed that region. This will be translated into the 3D velocity field as a blank in the data; no data will be shown for that region. This has been shown in the results of each individual view presented. However, the voids that appeared in the individual views are proven to be successfully filled in in the merged-view case. Thanks to the simplicity of probe positioning that Robotic PIV offers, there is an existing overlap between the measurement volumes, allowing the empty gaps in the data caused by SFFM to be compensated by another view. Furthermore, it has also been shown that when processing a high number of images (5,000 in this case), the spatial gaps introduced by the masking method are almost completely filled.

Regarding the propeller's velocity field results, both Minimum Subtraction and Butterworth present significant regions of outliers caused by the reflections from the struts and, more importantly, from the rotation of the blade. The latter could be observed in the YZ plane at the wake of the propeller. An important aspect is that SFFM showed low standard deviation values, indicating a good quality of the results.

While inspecting the results for the three geometries, there was something that was noticeable: the field of view of the SFFM was much smaller than the Butterworth's. This is because the Butterworth time filter's process includes an intensity normalization step that allows to have a uniform intensity across the whole image domain. By observing the SFFM images, one can notice that the particles further away from the centre of the image have a decaying intensity compared to the centre. Therefore, an intensity normalization step was performed at the end of the methodology pipeline. The resulting velocity field showed a significant increase of the field of view, without affecting the flow field results.

3D-based Particle Concentration Mask

Overall, this methodology failed at detecting and removing the regions of reflections over the three datasets from the experiments performed. It was applied on the Minimum Subtraction results with the aim to have a larger particle concentration separation between reflections and particles. However, as shown in the results, this separation was not sufficient for the method to effectively eliminate the reflection's contribution. The velocity field results provided a difficult to interpret data, that is far from the results of both state-of-the-art and the proposed image pre-processing techniques.

7.2. Recommendations

It has been shown that one of the three proposed methodologies shows great potential in reflection mitigation, yielding a good quality of the resulting 3D velocity field. However, there are a few opportunities for improvement to refine the study.

1. A set of 5,000 images were acquired for each view in DF-STB in the case of the side-view mirror and Formula 1 car and in TR-STB for the propeller. Due to time constraints in the project timeline, the post-processing was only done in 500 images per view. If a further accuracy is desired, the whole set of 5,000 images should be processed. This will make the data statistically more reliable, allowing to retrieve a larger amount of data over the measurement domain.
2. The original Spatial Fourier Filter or Spatial Fourier Filter + Mask processes did not consider the additional step of intensity normalization. However, it has been proven that it significantly extends the field of view obtained with these approaches. Hence, it is recommended to add this extra operation on the methodology pipeline with the aim to obtain a larger field of view, which is beneficial if a part of the flow field needs to be studied and lies within this extended FOV.
3. There was no reference results for neither of the geometries tested, hence the obtained results could not be completely validated. For future studies and by means of further validation, a test case should be selected and tested with two different systems: a bigger aperture system set-up and with Robotic PIV.
4. A possible future line of work that can further extend this project is to investigate a way of automated filling in the masked regions. The idea would be to create an algorithm that can predict the regions of the measurement volume that will be blank due to the SFFM implementation and, then propose new measurement volumes to acquire with Robotic PIV with the aim to have all the gaps filled in.
5. The 3D-based Particle Concentration Mask has been shown to not provide the desired results, not being able to identify the regions of reflection by particle concentration. However, this could be due to the way the number of particles is computed. A possible alternative to obtain this particle concentration could be defined, which could weight more the regions of reflections (regions with accumulation of particles) than regions with particles tracks.

References

[1] Zoheidi, L. (2018). *Flow Characterization of Milk Protein Foam Transport in Inner Geometries* (Doctoral dissertation). Technischen Fakultät der Friedrich-Alexander-Universität Erlangen-Nürnberg. https://www.researchgate.net/publication/327133859_Flow_Characterization_of_Milk_Protein_Foam_Transport_in_Inner_Geometries?_tp=eyJjb250ZXh0Ijp7ImZpcnN0UGFnZSI6Il9kaXJlY3QiLCJwYWdlIjoiX2RpcmVjdCJ9fQ

[2] LaVision. (n.d.[a]). Applications. https://www.lavision.de/en/applications/index.php

[3] Raffel, M., Willert, C. E., Scarano, F., Kähler, C. J., Wereley, S. T., & Kompenhans, J. (2018). Particle Image Velocimetry. https://doi.org/10.1007/978-3-319-68852-7

[4] Tropea, C., Yarin, A. L., & Foss, J. F. (2007). Springer Handbook of Experimental Fluid Mechanics. *Springer Handbook of Experimental Fluid Mechanics*. https://doi.org/10.1007/978-3-540-30299-5

[5] Meinhart, C. D., & Wereley, S. T. (2003). The theory of diffraction-limited resolution in microparticle image velocimetry. *Measurement Science and Technology*, *14*(7), 1047. https://doi.org/10.1088/0957-0233/14/7/320

[6] Harmsen, N. (2021). A Closer Look at Lens Diffraction | Fstoppers. https://fstoppers.com/science/closer-look-lens-diffraction-566322

[7] Hollows, G., & James, N. (n.d.). The Airy Disk and Diffraction Limit | Edmund Optics. https://www.edmundoptics.es/knowledge-center/application-notes/imaging/limitations-on-resolution-and-contrast-the-airy-disk/

[8] Willert, C. E., & Gharib, M. (1991). Digital particle image velocimetry. *Experiments in Fluids*, *10*(4), 181–193. https://doi.org/10.1007/BF00190388/METRICS

[9] Adrian, R. J., & Westerweel, J. (2010). Particle Image Velocimetry (Cambridge Aerospace Series). (30), 586. https://books.google.com/books/about/Particle_Image_Velocimetry.html?id=jbDl2-yHbooC

[10] Michaelis, D., Neal, D. R., & Wieneke, B. (2016). Peak-locking reduction for particle image velocimetry. *Measurement Science and Technology*, *27*(10), 104005. https://doi.org/10.1088/0957-0233/27/10/104005

[11] Kislaya, A. (2016). *Estimation and reduction of peak-locking errors in PIV measurements* (Doctoral dissertation). TU Delft. Delft.

[12] Giannopoulos, A., Passaggia, P. Y., Mazellier, N., & Aider, J. L. (2022). On the optimal window size in optical flow and cross-correlation in particle image velocimetry: application to turbulent flows. *Experiments in Fluids*, *63*(3). https://doi.org/10.1007/S00348-022-03410-Z

[13] LaVision. (n.d.[b]). Particle Image Velocimetry. https://www.lavision.de/en/techniques/piv-ptv/

[14] Guida, A., Nienow, A. W., & Barigou, M. (2010). The effects of the azimuthal position of the measurement plane on the flow parameters determined by PIV within a stirred vessel.

Chemical Engineering Science, *65*(8), 2454–2463. https://doi.org/10.1016/j.ces.2009.12.012

[15] Terra, W., Sciacchitano, A., & Scarano, F. (2016). Drag Analysis from PIV Data in Speed Sports. *Procedia Engineering*, *147*, 50–55. https://doi.org/10.1016/J.PROENG.2016.06.188

[16] Brito, M., Sanches, P., Ferreira, R. M. L., & Covas, D. I. C. (2017). Experimental study of the transient flow in a coiled pipe using PIV. *Journal of Hydraulic Engineering*, *143*(3), 04016087. https://doi.org/10.1061/(ASCE)HY.1943-7900.0001253

[17] Barbosa, F., Teixeira, S., Costa, C., Marques, F., & Teixeira, J. C. (2020). Experimental study of multiple air jets impinging a moving flat plate. *ASME International Mechanical Engineering Congress and Exposition, Proceedings (IMECE)*, *11*, V011T11A051. https://doi.org/10.1115/IMECE2020-23996

[18] Jux, C., Sciacchitano, A., & Scarano, F. (2023). Tire dependence for the aerodynamics of yawed bicycle wheels. *Journal of Wind Engineering and Industrial Aerodynamics*, *233*, 105294. https://doi.org/10.1016/j.jweia.2022.105294

[19] Arroyo, M. P., & Greated, C. A. (1991). Stereoscopic particle image velocimetry. *Meas. Sci. Technol.*, *2*(12), 1181–1186. https://doi.org/10.1088/0957-0233/2/12/012

[20] Spoelstra, A., de Martino Norante, L., Terra, W., Sciacchitano, A., & Scarano, F. (2019). On-site cycling drag analysis with the Ring of Fire. *Experiments in Fluids*, *60*(6), 90. https://doi.org/10.1007/s00348-019-2737-y

[21] Roloff, C., Stucht, D., Beuing, O., & Berg, P. (2019). Comparison of intracranial aneurysm flow quantification techniques: Standard PIV vs stereoscopic PIV vs tomographic PIV vs phase-contrast MRI vs CFD. *Journal of NeuroInterventional Surgery*, *11*(3), 275–282. https://doi.org/10.1136/neurintsurg-2018-013921

[22] Spoelstra, A., Sciacchitano, A., Scarano, F., & Mahalingesh, N. (2021). On-site drag analysis of drafting cyclists. *Journal of Wind Engineering and Industrial Aerodynamics*, *219*, 104797. https://doi.org/10.1016/j.jweia.2021.104797

[23] Huang, M., Sciacchitano, A., & Ferreira, C. (2023). On the wake deflection of vertical axis wind turbines by pitched blades. *Wind Energy*, *26*(4), 365–387. https://doi.org/10.1002/we.2803

[24] Huang, M., Vijaykumar Patil, Y., Sciacchitano, A., & Ferreira, C. (2023). Experimental study of the wake interaction between two vertical axis wind turbines. *Wind Energy*. https://doi.org/10.1002/we.2863

[25] Matt, S., Nootz, G., Hellman, S., & Hou, W. (2020). Effects of Optical Turbulence and Density Gradients on Particle Image Velocimetry. *Scientific Reports 2020 10:1*, *10*(1), 1–12. https://doi.org/10.1038/s41598-020-58077-5

[26] Kähler, C. J. (2004). Investigation of the spatio-temporal flow structure in the buffer region of a turbulent boundary layer by means of multiplane stereo PIV. *Experiments in Fluids*, *36*(1), 114–130. https://doi.org/https://doi.org/10.1007/s00348-003-0680-3

[27] Kähler, C. J. (2006). The significance of turbulent eddies for the mixing in boundary layers. *Solid Mechanics and its Applications*, *129*, 405–414. https://doi.org/https://doi.org/10.1007/978-1-4020-4150-1{_}39

[28] Schröder, F., Klaas, M., & Schröder, W. (2011). Multiplane-Stereo PIV measurements for steady flow in the first two bifurcations of the upper human airways during exhalation. *Conference on Simulation and Modeling of Biological Flows*.

[29] Elsinga, G. E., Scarano, F., Wieneke, B., & Van Oudheusden, B. W. (2006). Tomographic particle image velocimetry. *Experiments in Fluids*, *41*(6), 933–947. https://doi.org/10.1007/s00348-006-0212-z

[30] Scarano, F., Ghaemi, S., Caridi, G. C. A., Bosbach, J., Dierksheide, U., & Sciacchitano, A. (2015). On the use of helium-filled soap bubbles for large-scale tomographic PIV in wind tunnel experiments. *Experiments in Fluids*, *56*(2), 1–12. https://doi.org/https://doi.org/10.1007/s00348-015-1909-7

[31] Bosbach, J., Kühn, M., & Wagner, C. (2009). Large scale particle image velocimetry with helium filled soap bubbles. *Experiments in Fluids*, *46*(3), 539–547. https://doi.org/https://doi.org/10.1007/s00348-008-0579-0

[32] Zang, W., & Prasad, A. K. (1997). Performance evaluation of a Scheimpflug stereocamera for particle image velocimetry. *Applied Optics, Vol. 36, Issue 33, pp. 8738-8744*, *36*(33), 8738–8744. https://doi.org/10.1364/AO.36.008738

[33] Atkinson, C., & Soria, J. (2009). An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Experiments in Fluids*, *47*(4-5), 553–568. https://doi.org/10.1007/S00348-009-0728-0/METRICS

[34] LaVision. (n.d.[c]). Tomographic PIV. https://www.lavision.de/en/products/flowmaster/tomographic-piv/

[35] Wieneke, B. (2008). Volume self-calibration for 3D particle image velocimetry. *Experiments in Fluids*, *45*(4), 549–556. https://doi.org/10.1007/S00348-008-0521-5/TABLES/1

[36] Schanz, D., Gesemann, S., & Schröder, A. (2016). Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Experiments in Fluids*, *57*(5), 1–27. https://doi.org/https://doi.org/10.1007/s00348-016-2157-1

[37] Jahn, T., Schanz, D., & Schröder, A. (2021). Advanced iterative particle reconstruction for Lagrangian particle tracking. *Experiments in Fluids*, *62*(8), 1–24. https://doi.org/https://doi.org/10.1007/s00348-021-03276-7

[38] Saredi, E., Sciacchitano, A., & Scarano, F. (2020). Multi-Δt 3D-PTV based on Reynolds decomposition. *Measurement Science and Technology*, *31*(8), 084005. https://doi.org/10.1088/1361-6501/ab803d

[39] Jux, C., Sciacchitano, A., Schneiders, J. F., & Scarano, F. (2018). Robotic volumetric PIV of a full-scale cyclist. *Experiments in Fluids*, *59*(4), 1–15. https://doi.org/https://doi.org/10.1007/s00348-018-2524-1

[40] Schneiders, J. F., Scarano, F., Jux, C., & Sciacchitano, A. (2018). Coaxial volumetric velocimetry. *Measurement Science and Technology*, *29*(6), 065201. https://doi.org/10.1088/1361-6501/AAB07D

[41] Xue, Z., Charonko, J. J., & Vlachos, P. P. (2014). Particle image velocimetry correlation signal-to-noise ratio metrics and measurement uncertainty quantification. *Measurement Science and Technology*, *25*(11), 115301. https://doi.org/10.1088/0957-0233/25/11/115301

[42] Hannemann, K., Lüdeke, H., Pallegoix, J., Ollivier, A., Lambaré, H., Maseland, H., Geurts, E., Frey, M., Deck, S., Schrijer, F., Scarano, F., & Schwane, R. (2011). Launch Vehicle Base

Buffeting - Recent Experimental and Numerical Investigations. *7th European Symposium on Aerothermodynamics*.

[43] Cadel, D. R., Shin, D., & Todd Lowe, K. (2016). A hybrid technique for laser flare reduction. *54th AIAA Aerospace Sciences Meeting*, *0*. https://doi.org/10.2514/6.2016-0788

[44] Schuster, U., Quosdorf, D., Ritter, K., & Manhart, M. (2015). Surface reflection of 3D scour models in Particle-Image-Velocimetry Experiments.

[45] Paterna, E., Moonen, P., Dorer, V., & Carmeliet, J. (2013). Mitigation of surface reflection in PIV measurements. *Measurement Science and Technology*, *24*(5), 057003. https://doi.org/10.1088/0957-0233/24/5/057003

[46] Petrosky, B. J., Lowe, K. T., Danehy, P. M., Wohl, C. J., & Tiemsin, P. I. (2015). Improvements in laser flare removal for particle image velocimetry using fluorescent dye-doped particles. *Measurement Science and Technology*, *26*(11), 115303. https://doi.org/10.1088/0957-0233/26/11/115303

[47] Petrosky, B. J., Maisto, P., Todd Lowe, K., André, M. A., Bardet, P. M., Tiemsin, P. I., Wohl, C. J., & Danehy, P. M. (2015). Particle image velocimetry applications using fluorescent dye-doped particles. *53rd AIAA Aerospace Sciences Meeting*. https://doi.org/10.2514/6.2015-1223

[48] Acharya, A. S., Lowe, K. T., & Ng, W. F. (2022). Fluorescent particle image velocimetry using atomized liquid particles. *Measurement Science and Technology*, *33*(6), 065301. https://doi.org/10.1088/1361-6501/AC543B

[49] Lin, H. J., & Perlin, M. (1998). Improved methods for thin, surface boundary layer investigations. *Experiments in Fluids*, *25*(5-6), 431–444. https://doi.org/10.1007/S003480050249/METRICS

[50] Kähler, C. J., Scholz, U., & Ortmanns, J. (2006). Wall-shear-stress and near-wall turbulence measurements up to single pixel resolution by means of long-distance micro-PIV. *Experiments in Fluids*, *41*(2), 327–341. https://doi.org/10.1007/S00348-006-0167-0/FIGURES/17

[51] Westerweel, J. (1994). Efficient detection of spurious vectors in particle image velocimetry data sets. *Exp Fluids*, *16*(3-4), 236–247. https://doi.org/10.1007/bf00206543

[52] Westerweel, J., & Scarano, F. (2005). Universal outlier detection for PIV data. *Experiments in Fluids*, *39*(6), 1096–1100. https://doi.org/10.1007/S00348-005-0016-6/FIGURES/5

[53] Wang, H. P., Gao, Q., Feng, L. H., Wei, R. J., & Wang, J. J. (2015). Proper orthogonal decomposition based outlier correction for PIV data. *Experiments in Fluids*, *56*(2), 1–15. https://doi.org/10.1007/S00348-015-1894-X

[54] Higham, J. E., Brevis, W., & Keylock, C. J. (2016). A rapid non-iterative proper orthogonal decomposition based outlier detection and correction for PIV data. *Measurement Science and Technology*, *27*(12). https://doi.org/10.1088/0957-0233/27/12/125303

[55] Tang, C., Sun, W., He, H., Li, H., & Li, E. (2017). Spurious PIV vector detection and correction using a penalized least-squares method with adaptive order differentials. *Experiments in Fluids*, *58*(7). https://doi.org/10.1007/S00348-017-2350-X

[56] Saredi, E., Sciacchitano, A., & Scarano, F. (2022). Outlier detection for PIV statistics based on turbulence transport. *Experiments in Fluids*, *63*(1). https://doi.org/10.1007/S00348-021-03368-4

[57] Gui, L., & Merzkirch, W. (2001). A comparative study of the MQD method and several correlation-based PIV evaluation algorithms. *Experiments in Fluids*, *28*(1), 36–44. https://doi.org/10.1007/S003480050005/METRICS

[58] Wereley, S. T., Gui, L., & Meinhart, C. D. (2012). Advanced Algorithms for Microscale Particle Image Velocimetry. *https://doi.org/10.2514/2.1786*, *40*(6), 1047–1055. https://doi.org/10.2514/2.1786

[59] LaVision. (n.d.[d]). PTV / Shake-the-Box. https://www.lavision.de/en/products/flowmaster/3d-ptv-shake-the-box/index.php

[60] Honkanen, M., & Nobach, H. (2005). Background extraction from double-frame PIV images. *Exp. Fluids*, *38*(3), 348–362. https://doi.org/10.1007/s00348-004-0916-x

[61] Mejia-Alvarez, R., & Christensen, K. T. (2013). Robust suppression of background reflections in PIV images. *Meas. Sci. Technol.*, *24*(2), 027003. https://doi.org/10.1088/0957-0233/24/2/027003

[62] Lee, Y., Yang, H., & Yin, Z. (2017). Outlier detection for particle image velocimetry data using a locally estimated noise variance. *Measurement Science and Technology*, *28*(3). https://doi.org/10.1088/1361-6501/AA5431

[63] Wang, L., Pan, C., Liu, J., & Cai, C. (2020). Ratio-cut background removal method and its application in near-wall PTV measurement of a turbulent boundary layer. *Measurement Science and Technology*, *32*(2), 025302. https://doi.org/10.1088/1361-6501/ABB483

[64] Sciacchitano, A., & Scarano, F. (2014). Elimination of PIV light reflections via a temporal high pass filter. *Measurement Science and Technology*, *25*(8), 084009. https://doi.org/10.1088/0957-0233/25/8/084009

[65] Mendez, M. A., Raiola, M., Masullo, A., Discetti, S., Ianiro, A., Theunissen, R., & Buchlin, J. M. (2017). POD-based background removal for particle image velocimetry. *Exp. Therm. Fluid Sci.*, *80*, 181–192. https://doi.org/10.1016/j.expthermflusci.2016.08.021

[66] Adatrao, S., & Sciacchitano, A. (2019). Elimination of unsteady background reflections in PIV images by anisotropic diffusion. *Measurement Science and Technology*, *30*(3), 035204. https://doi.org/10.1088/1361-6501/AAFCA9

[67] Gonzalez, R. C., & Woods, R. E. (2018). *Digital Image Processing* (Fourth). Pearson/Prentice Hall.

[68] Dictionary:Fourier transform - SEG Wiki. (n.d.). https://wiki.seg.org/wiki/Dictionary:Fourier_transform

[69] Arduino Tutorial: Simple High-pass, Band-pass and Band-stop Filtering – Norwegian Creations. (n.d.). https://www.norwegiancreations.com/2016/03/arduino-tutorial-simple-high-pass-band-pass-and-band-stop-filtering/

[70] Gibbs Phenomenon – from Wolfram MathWorld. (n.d.). https://mathworld.wolfram.com/GibbsPhenomenon.html

[71] Nixon, M. S., & Aguado, A. S. (2019). *Feature extraction and image processing for computer vision*. Elsevier. https://doi.org/10.1016/C2017-0-02153-5

[72] Understanding 2D Fourier space. Build your 2D Fourier intuitions for… | by Peter Barrett Bryan | Medium. (n.d.). https://peterbbryan.medium.com/understanding-2d-fourier-space-59808b644a13

[73] Kundur, D. (2013). University of Toronto Notes. https://ww2.comm.utoronto.ca/~dkundur//course_info/signals/notes/

[74] Shechtman, Y., Eldar, Y. C., Cohen, O., Chapman, H. N., Miao, J., & Segev, M. (2015). Phase Retrieval with Application to Optical Imaging: A contemporary overview. *IEEE Signal Processing Magazine*, *32*(3), 87–109. https://doi.org/10.1109/MSP.2014.2352673

[75] Gallagher, T. A., Nemeth, A. J., & Hacein-Bey, L. (2008). An introduction to the Fourier transform: Relationship to MRI. *American Journal of Roentgenology*, *190*(5), 1396–1405. https://doi.org/10.2214/AJR.07.2874/ASSET/IMAGES/05{_}07{_}2874{_}13B{_}CMYK.JPEG

[76] Fahimian, B. P., Zhao, Y., Huang, Z., Fung, R., Mao, Y., Zhu, C., Khatonabadi, M., Demarco, J. J., Osher, S. J., McNitt-Gray, M. F., & Miao, J. (2013). Radiation dose reduction in medical x-ray CT via Fourier-based iterative reconstruction. *Medical Physics*, *40*(3), 31914–31915. https://doi.org/10.1118/1.4791644

[77] Hansen, M. S., & Kellman, P. (2015). Image reconstruction: An overview for clinicians. *Journal of Magnetic Resonance Imaging*, *41*(3), 573–585. https://doi.org/10.1002/JMRI.24687

[78] Lai, J. H., Yuen, P. C., & Feng, G. C. (1999). Spectroface: A Fourier-based approach for human face recognition. https://scholars.hkbu.edu.hk/en/publications/spectroface-a-fourier-based-approach-for-human-face-recognition

[79] Samra, A. S., El Taweel Gad Allah, S., & Ibrahim, R. M. (2003). Face Recognition Using Wavelet Transform, Fast Fourier Transform and Discrete Cosine Transform. *Midwest Symposium on Circuits and Systems*, *1*, 272–275. https://doi.org/10.1109/MWSCAS.2003.1562271

[80] Gouru, A., & Suthaharan, S. (2021). Facial Emotion Characterization and Detection using Fourier Transform and Machine Learning. *EPiC Series in Computing*, *82*, 150–159. https://doi.org/10.29007/7kbw

[81] Hu, J., Deng, J., & Wu, J. (2011). Image compression based on improved FFT algorithm. *Journal of Networks*, *6*(7), 1041–1048. https://doi.org/10.4304/JNW.6.7.1041-1048

[82] Sahnoun, K., & Benabadji, N. (2013). On-board satellite image compression using the Fourier transform and Huffman coding. *2013 World Congress on Computer and Information Technology, WCCIT 2013*. https://doi.org/10.1109/WCCIT.2013.6618754

[83] Bretschneider, T. (2002). On the deconvolution of satellite imagery. *International Geoscience and Remote Sensing Symposium (IGARSS)*, *4*, 2450–2452. https://doi.org/10.1109/IGARSS.2002.1026574

[84] Kesh, S., & Chakravarty, S. (2011). Application of fourier transform to get an original satellite image without applying the SRM or 180 degree rotation. *Communications in Computer and Information Science*, *260 CCIS*, 48–55. https://doi.org/10.1007/978-3-642-27183-0{_}6/COVER

[85] Fan, Y., Guo, C., Han, Y., Qiao, W., Xu, P., & Kuai, Y. (2023). Deep-learning-based image preprocessing for particle image velocimetry. *Applied Ocean Research*, *130*, 103406. https://doi.org/10.1016/J.APOR.2022.103406

[86] Jux, C., Sciacchitano, A., & Scarano, F. (2020). Flow pressure evaluation on generic surfaces by robotic volumetric PTV. *Measurement Science and Technology*, *31*(10), 104001. https://doi.org/10.1088/1361-6501/AB8F46

[87] Westerweel, J. (1997). Fundamentals of digital particle image velocimetry. *Measurement Science and Technology*, *8*(12), 1379. https://doi.org/10.1088/0957-0233/8/12/002

[88] Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, *13*(4), 600–612. https://doi.org/10.1109/TIP.2003.819861

[89] Caridi, G. C. A., Ragni, D., Sciacchitano, A., & Scarano, F. (2016). HFSB-seeding for large-scale tomographic PIV in wind tunnels. *Experiments in Fluids*, *57*(12), 1–13. https://doi.org/10.1007/S00348-016-2277-7/FIGURES/13

[90] Open Jet Facility. (n.d.). https://www.tudelft.nl/lr/organisatie/afdelingen/flow-physics-and-technology/facilities/low-speed-wind-tunnels/open-jet-facility

[91] Gardiner, J. D., Behnsen, J., & Brassey, C. A. (2018). Alpha shapes: Determining 3D shape complexity across morphologically diverse structures. *BMC Evolutionary Biology*, *18*(1), 1–16. https://doi.org/10.1186/S12862-018-1305-Z/FIGURES/10

[92] Tu Delft. (n.d.). W-Tunnel. https://www.tudelft.nl/lr/organisatie/afdelingen/flow-physics-and-technology/facilities/low-speed-wind-tunnels/w-tunnel

[93] Saredi, E., Tumuluru Ramesh, N., Sciacchitano, A., & Scarano, F. (2021). State observer data assimilation for RANS with time-averaged 3D-PIV data. *Computers & Fluids*, *218*, 104827. https://doi.org/10.1016/J.COMPFLUID.2020.104827

[94] Cueto Corral, P. (2023). *Propeller Slipstream Characterisation by Large-Scale PIV: An experimental study for Isolated Propeller and Propeller-Wing Interaction* (Doctoral dissertation). TU Delft. Delft. https://repository.tudelft.nl/islandora/object/uuid%3A6d071306-faef-4cb9-9c21-a184fc8aaf0e

[95] RoboDK Getting Started UR5 Paint Gun Tool - YouTube. (n.d.). https://www.youtube.com/watch?app=desktop&v=CGUvYc5_Y78

[96] UR5e Lightweight, versatile cobot. (n.d.). https://www.universal-robots.com/products/ur5-robot/

[97] Grille Guerra, A., Scarano, F., & Sciacchitano, A. (2024). On the scalability of helium-filled soap bubbles for volumetric PIV. *Experiments in Fluids*, *65*(2). https://doi.org/10.1007/S00348-024-03760-W

[98] Agüera, N., Cafiero, G., Astarita, T., & Discetti, S. (2016). Ensemble 3D PTV for high resolution turbulent statistics. *Measurement Science and Technology*, *27*(12), 124011. https://doi.org/10.1088/0957-0233/27/12/124011

A Algorithms

This appendix contains the algorithms developed for the proposed reflection treatment methodologies.

A.1. Spatial Fourier filter method

A.1.1. Implementation of Fourier high-pass filter

```
function [filtered_image] = FT_Filter(input_image, kappaf)
    % Inputs:
    %    input_image - Input image
    %    kappaf - Filter's cut-on wavenumber/standard deviation [1/
       pixel]
    %
    % Outputs:
    %    filtered_image - Filtered image

    % Get Fourier Transform of the input_image
    % using fft2 (2D fast fourier transform)
    FT_img = fftshift(fft2(input_image));

    % Generate Gaussian High-pass filter with standard deviation =
       kappaf
    filter = generate_filter(I, J, kappaf);

    % Mask the Fourier Transform by multiplying by the filter
    G = FT_img.*filter;

    % Get filtered image in the space domain
    % by using the Inverse Fourier Transform operation
    filtered_image = uint16(ifft2(ifftshift(G),'symmetric'));

    if isreal(filtered_image) == false
        fprintf('WARNING: IFFT is complex \n')
    end

    filtered_image(filtered_image<0) = 0;
```

```
28 end
```

A.1.2. Generation of Gaussian high-pass filter

```
function [filter] = generate_filter(I, J, kappaf)
    % Inputs:
    %    I - Image width [pixels]
    %    J - Image height [pixels]
    %    kappaf - Filter's cut-on wavenumber/standard deviation [1/
       pixel]
    %
    % Outputs:
    %    filter - Gaussian high-pass filter

    % Define the x and y grid in wavenumber space
    % (Nyquist frequency is 0.5 cycles/px)
    x = linspace(-0.5,0.5-1/I,I);
    y = linspace(-0.5,0.5-1/J,J);
    [X,Y] = meshgrid(x,y);

    % Generate Gaussian high-pass filter
    filter = 1 - exp(-X.^2/(2*kappaf^2)).* exp(-Y.^2/(2*kappaf^2));
end
```

A.2. Spatial Fourier filter and mask method

```
function [masked_image] = Fourier_filter_mask_implementation(
   input_image, filtered_image, SSIM_threshold, Area_threshold)
    % Inputs:
    %    input_image - Input image
    %    filtered_image - Fourier-filtered image
    %    SSIM_threshold - SSIM value threshold [-]
    %    Area_threshold - Area value threshold [pixel^2]
    %
    % Outputs:
    %    masked_image - Masked image

    % Obtain SSIM map from comparing input and filtered images
    [ssimval,ssimmap] = ssim(filtered_image, input_image);

    % Obtain mask from thresholding SSIM coefficient plot
    ssimmapmask = ssimmap;
    ssimmapmask(ssimmapmask>SSIM_threshold) = 1;
    ssimmapmask(ssimmapmask<=SSIM_threshold) = 0;

    % Retrieve contours from Structural similarity (SSIM) plot
    contours = ssimmapmask == 0;
    [L, num] = bwlabel(contours, 8);

    % Count number of pixels in each blob (area of each blob)
    counts = sum(bsxfun(@eq,L(:),1:num));
```

```
25
26      % Find contours biggest than a set threshold area
27      [~,ind] = find(counts > Area_threshold);
28
29      % Generate mask with largest contours
30      num_indxs = size(ind,2);
31      if num_indxs > 0
32          biggestContours = ones(J, I);
33          for i = 1:num_indxs
34              K = L;
35              out = (K==ind(i));
36              if i == 1
37                  biggestContours = out & biggestContours;
38              else
39                  biggestContours = out + biggestContours;
40              end
41          end
42          mask = uint16(~biggestContours);
43      end
44
45      % Mask image
46      masked_image = filtered_image.*mask;
47  end
```

B

Side-view mirror

B.1. Instantaneous images for each view

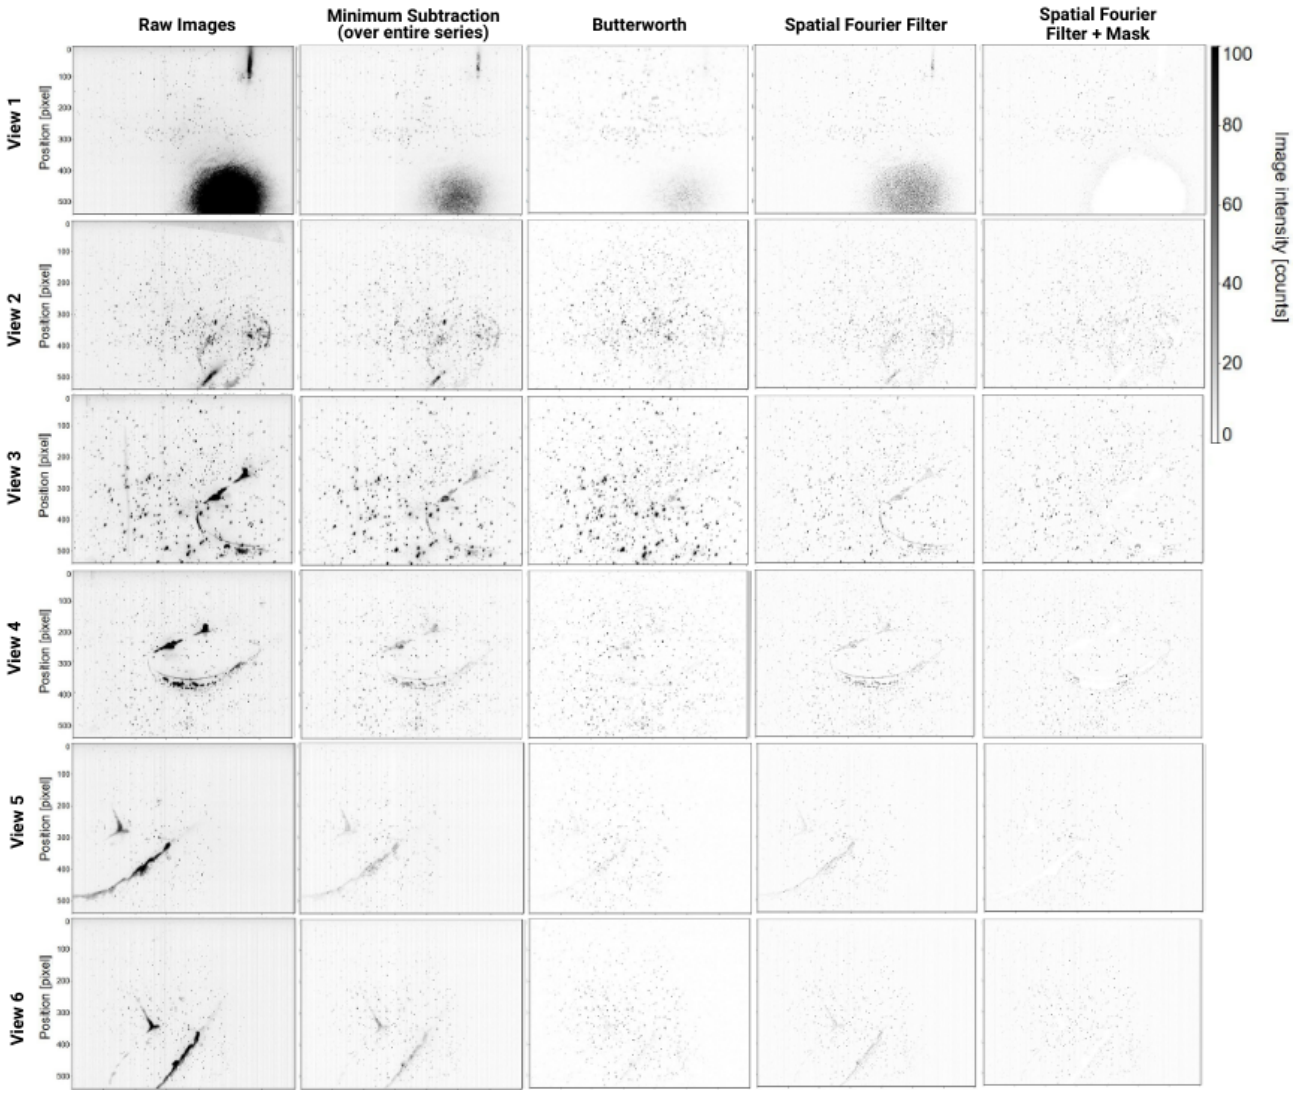

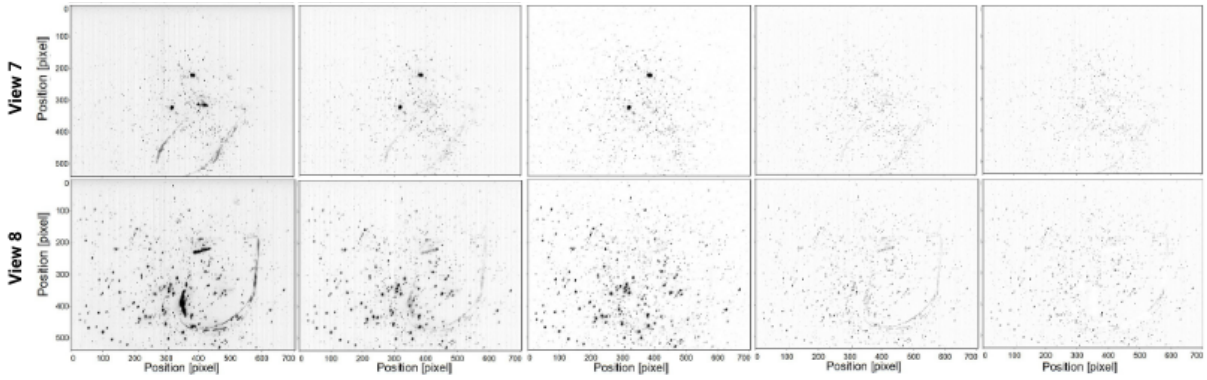

Figure B.1: Side-view mirror raw and pre-processed images from Views 1-8.

C

Formula 1 Car

C.1. Instantaneous images for each view

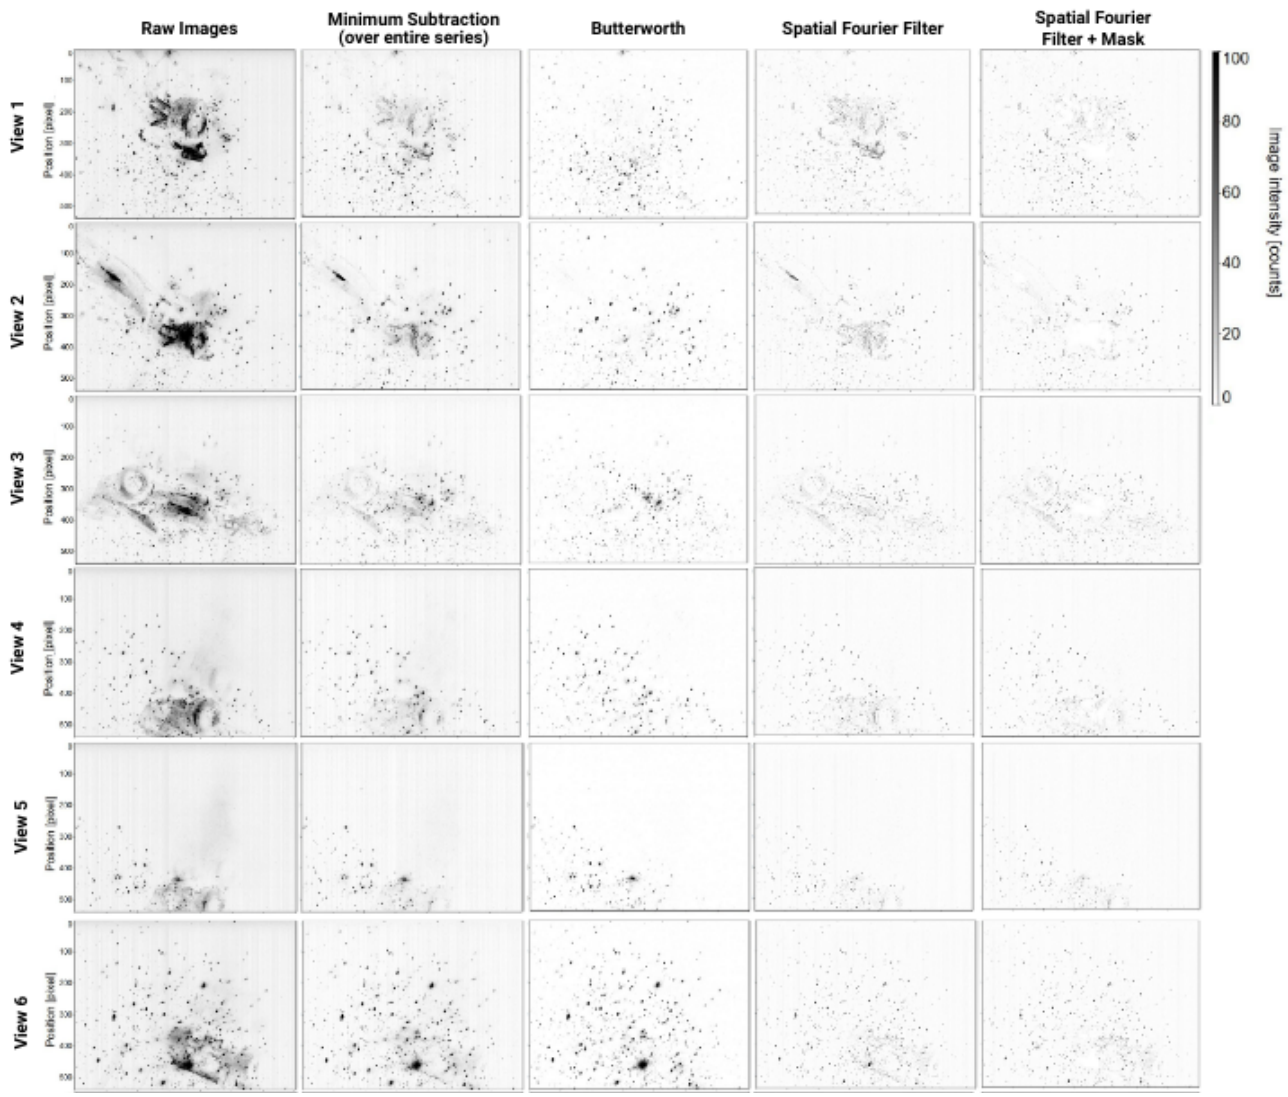

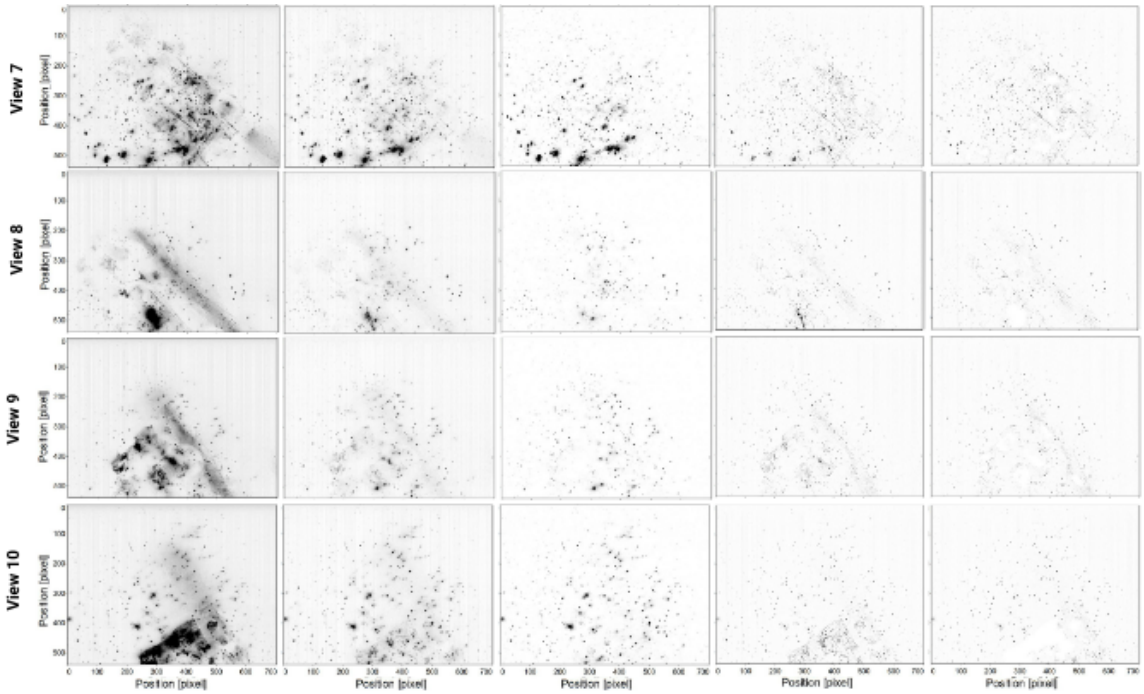

Figure C.1: Formula 1 car raw and pre-processed images from Views 1-10.

D

Propeller

D.1. Instantaneous images for each view

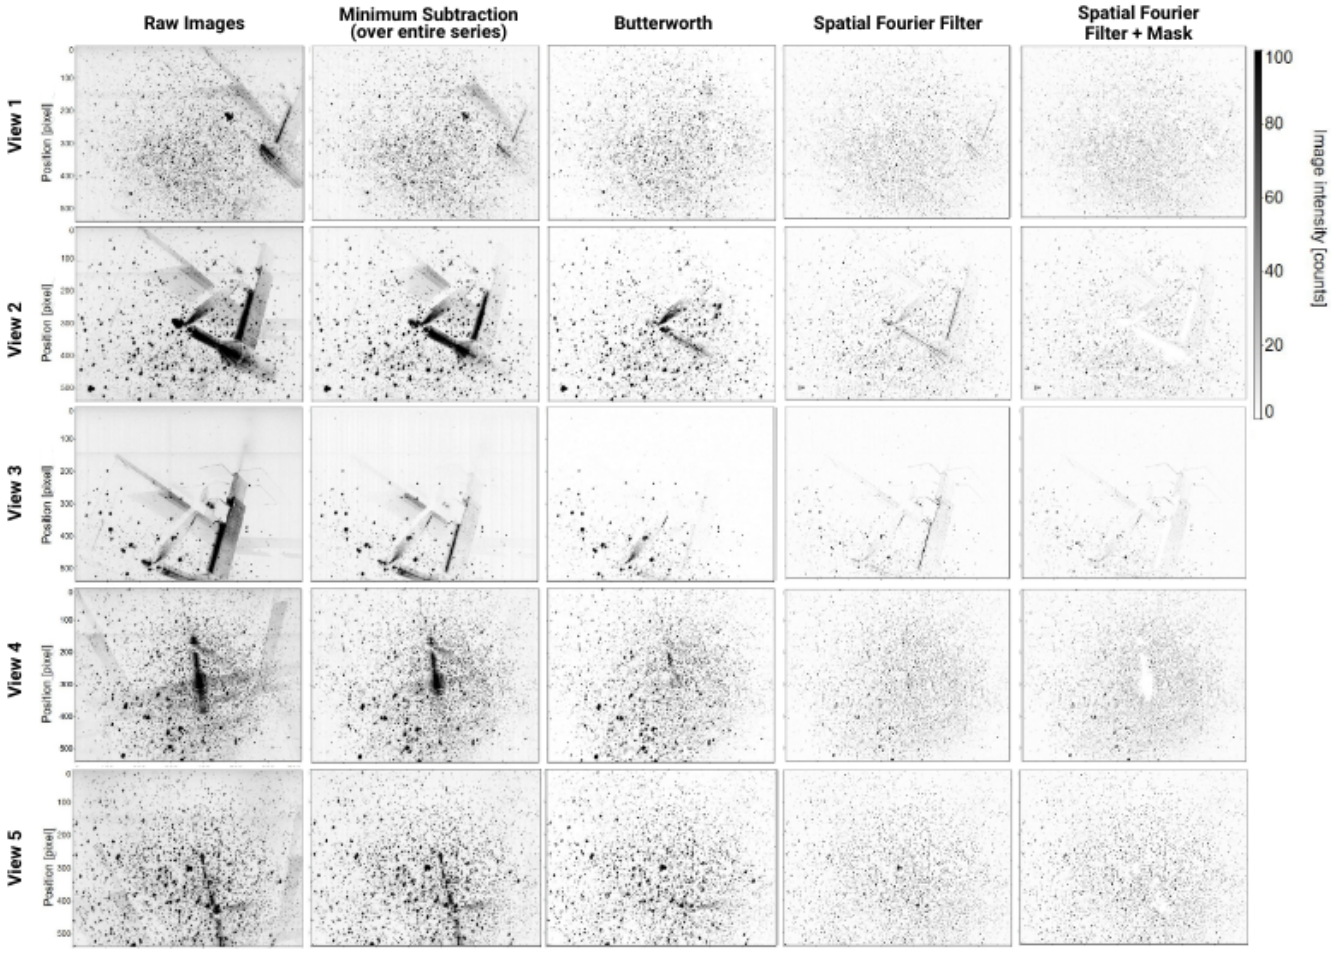

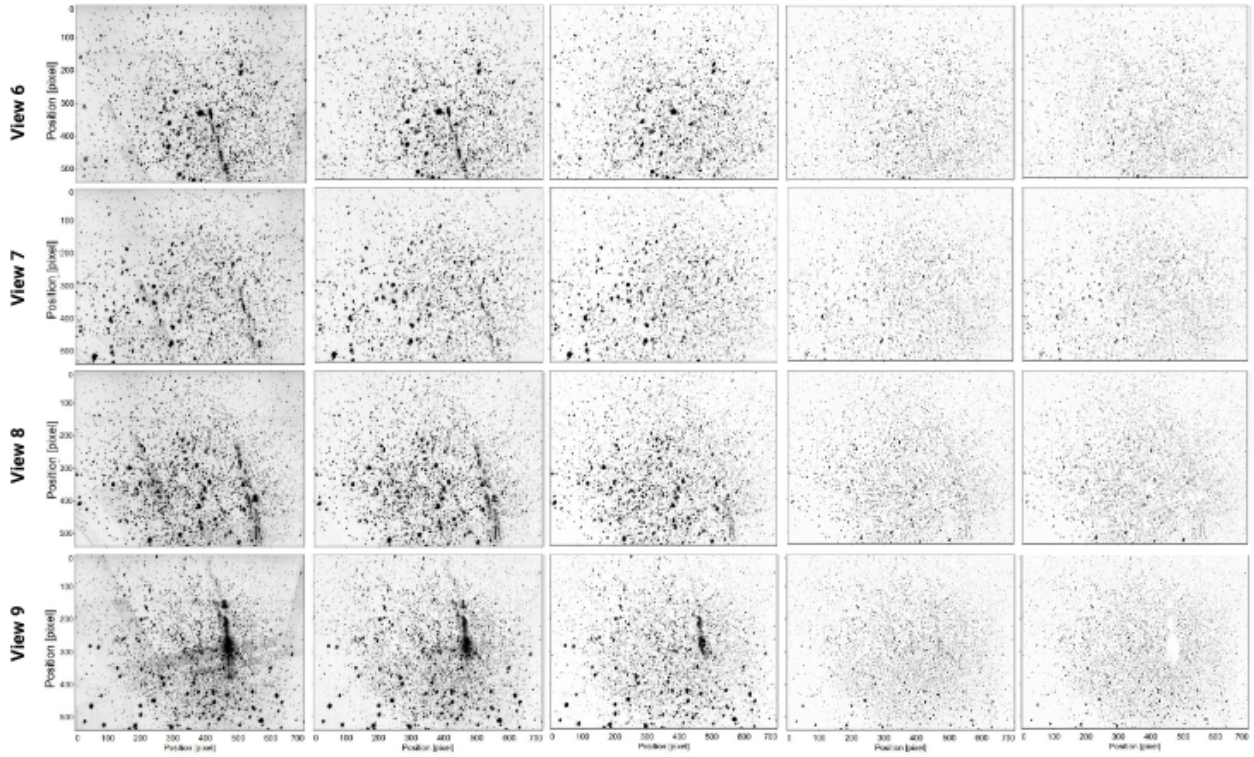

Figure D.1: Propeller raw and pre-processed images from Views 1-9.