

In-flight experimental measurements of large-scale deformations on a leading-edge inflatable kite

using stereoscopic photogrammetry and UWB modules

P. J. Haanen



Cover image: The V3 kite in flight at the Kitepower B.V. test site in Bangor Erris, Ireland.

DELFT UNIVERSITY OF TECHNOLOGY
FACULTY OF AEROSPACE ENGINEERING
MSC THESIS - AE5122

In-flight experimental measurements of large-scale deformations on a leading-edge inflatable kite

using stereoscopic photogrammetry and UWB modules

Instructors: Dr.-Ing. R. Schmehl, ir. J.A.W. Poland, Dr. C. Vergara
Pim Julius Haanen, 5092795
January 8, 2026

To obtain the degree of Master of Science in Aerospace Engineering at Delft University of Technology. To be defended publicly on 15th of January, 2026 at 10.00.

Student number: 5092795
Thesis duration: February 10, 2025 – January 15, 2026
Thesis committee: Dr.ir. D. A. M. De Tavernier TU Delft, chair
Dr. A. Sciacchitano TU Delft, examiner
Dr.-Ing. R. Schmehl TU Delft, supervisor
ir. J. A. W. Poland TU Delft, supervisor

An electronic version of this thesis is available at <https://repository.tudelft.nl/>.



This project was supported by the MERIDIONAL project, which has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101084216.

Preface

My journey with Kitepower B.V. began during the Design Synthesis Project at TU Delft, where I was first introduced to the concept of airborne wind energy (AWE). This initial interest was strengthened by a course on airborne wind energy led by Professor Roland Schmehl, whose insights into the technology inspired me to explore this field in greater depth. During this period, I decided to pursue an internship at Kitepower in order to engage more closely with this innovative technology. There, I gained a deeper understanding of the system and its potential role in the future renewable energy mix. Building on this experience, I embarked on the present MSc thesis project, carried out from 10 February 2025 to 15 January 2026.

Regarding the thesis, I would like to express my sincere gratitude to my daily supervisors, Professor R. Schmehl and PhD candidate J. A. W. (Jelle) Poland. This thesis would not have been possible without their excellent supervision, continuous availability for questions, constructive feedback, and willingness to brainstorm solutions whenever I encountered difficulties. I am also grateful to Claudio Vergara for his support as company supervisor.

Furthermore, I would like to thank PhD candidate O. (Oriol) Cayon for his continuous support and involvement, as well as Professor R. (Reza) Sabzevari for his assistance with computer-vision-related questions.

On top of this, I am thankful to several people at Kitepower. In particular, I would like to thank Wolf Fisscher, Yoshi Verspaget and Tom Lolies, all members of the hardware team, for their help with hardware, 3D printing and many other practical aspects that were essential for this thesis. Additionally, I wish to thank Akshay Raju and Théophile Dubois for their help with the laborious task of applying a large number of cross markers to the kite. A special word of thanks goes to Joren Bredael for assisting in the wind-tunnel tests and for collaborating on the optimisation of the vane design using the theoretical insights developed in this work. I am also grateful to Louis Taillandier for his assistance during the first wind-tunnel test campaign.

I would also like to thank my fellow students who were working at Kitepower during the course of this thesis: Daan van Wolffelaar, Bart Kroese and Andrea Bertozzi. Their continuous discussions, feedback and shared experiences contributed significantly to the progress of our respective projects and to the enjoyable working atmosphere at the company.

Finally, I would like to express my gratitude to the operations team in Bangor Erris, Ireland, where the experimental campaign was carried out. Together with Wolf Fisscher and Tom Lolies, they provided invaluable support and advice throughout the test week, ensuring that the experiments could be performed successfully; Mark Devas, Padraic Doherty, Stijn Gielen and Pepijn Ubbels. Without the help and dedication of all the people mentioned above, the results presented in this thesis would not have been possible.

*P.J. Haanen
Delft, December 2025*



Figure 1: Me flying a kite on the beach (Texel, The Netherlands, 2010).

Abstract

Airborne wind energy (AWE) systems harness strong and persistent winds at altitudes beyond the reach of conventional turbines, offering an alternative for renewable power generation. Among the various concepts, the pumping-cycle soft-wing kite system has demonstrated significant potential, in which a tethered, flexible wing flies crosswind trajectories to generate power during reel-out and retraction phases. However, the performance and controllability of such systems depend strongly on the aerodynamic behaviour of the deformable wing, which undergoes large-scale structural deflections under load. To ensure safe operation and to improve aero-structural models used in simulation-based design, in-flight quantification of these deformations is essential.

This study presents the first successful in-flight measurement of large-scale deformations of a soft-wing kite using a fully integrated stereoscopic photogrammetry and ultra-wideband (UWB) ranging setup. The experiments were conducted at Kitepower’s test site in Bangor Erris, Ireland, using the V3 kite equipped with the newly developed Γ -bar—a rigid glass-fibre cross-bar integrated within the bridle system. Two synchronised wide-lens GoPro cameras and a Pitot tube were mounted on the bar to capture the three-dimensional (3D) wing geometry and apparent wind speed throughout flight. UWB tags placed near the wing tips provided independent span estimates for sensor fusion. All subsystems were synchronised to obtain time-resolved deformation data during representative flight manoeuvres, including steady flight, depowering, and turning.

A custom shape reconstruction pipeline was developed to process the stereoscopic video recordings. The algorithm performs camera calibration, feature detection, triangulation, and temporal tracking to recover the wing’s 3D surface in flight. The resulting spatial accuracy was within centimetre range across the entire wingspan (accuracy: 1.22% of the span), sufficient to resolve global deformation modes such as spanwise contraction, and wing shear and twist. The UWB data were used to correct for bending and torsional deflections of the Γ -bar itself. Through a sensor-fusion step in post-processing, the UWB-derived span provides an adjustive calibration that compensates for bar bending and twist, thereby improving the global consistency of the photogrammetric throughout the manoeuvres.

The results demonstrate that the kite undergoes distinct and repeatable large-scale deformations in flight. Relative to the undeformed CAD geometry, the in-flight wing shows pronounced billowing across the canopy, contracting the TE, indicating that future aerodynamic models should incorporate this contracted and cambered shape rather than relying on the nominal CAD surface. During depowering, the wing transitions to a more curved, anhedral-increased configuration: the effective angle of attack is reduced, the tether force drops substantially, and the wingspan decreases. In turning manoeuvres, a clear asymmetry develops between the pulled and released side. The pulled side experiences a higher effective angle of attack and a locally increased projected area, while the opposite occurs on the released side. This results in a rearward displacement of the pulled tip, producing bottom- and front-view shear, resulting in force asymmetries that induce rotational moments. Static validation tests were also performed to provide independent reference cases for structural model development. The suspended-kite measurements confirmed that the experimental setup can capture relevant deformation patterns, and that the simplified structural model reproduces the qualitative behaviour under controlled loads, though modelling accuracy could be improved.

This research establishes a robust, field-validated framework for in-flight deformation measurement of flexible membrane wings under realistic operational conditions. The method provides experimental data for validating coupled aero-structural models and opens the path toward real-time monitoring and control. Ultimately, the developed measurement pipeline contributes to safer operation, improved energy yield, and accelerated design optimisation of next-generation airborne wind energy systems.

Contents

Preface	i
Abstract	ii
List of Figures	vi
List of Tables	xi
1 Introduction	1
1.1 Airborne wind energy & the pumping cycle	1
1.2 The V3 kite	3
1.3 Motivation	3
1.4 Report layout	4
2 Literature review	6
2.1 Aerodynamics of a LEI kite	6
2.2 Modelling a kite	8
2.3 Typical deformations	10
2.3.1 Depowering deformation	11
2.3.2 Steering-induced deformations	11
2.3.3 Vibrational deformations	12
2.4 Previous experimental kite studies	13
2.4.1 Wind-tunnel campaigns and controlled experiments	13
2.4.2 In-flight visual tracking	14
2.4.3 Distributed kinematic/structural sensing on the wing	15
2.4.4 Wind-speed measurements	16
2.5 Measurement techniques investigated for kite deformation measurements	17
2.5.1 Photogrammetry	17
2.5.2 Ultra-wideband modules (UWB)	19
2.5.3 Others	19
2.6 Summary and outlook	21
3 Conceptual research design	22
3.1 Research objective	22
3.2 Research questions	22
3.3 Selection of measurement techniques	23
4 Stereoscopic photogrammetry setup	24
4.1 Basic principle	24
4.2 GoPro Hero9 Black	25
4.3 Camera calibration	26
4.3.1 Single-camera calibration	27
4.3.2 Stereoscopic calibration	31
4.4 Markers	35
4.4.1 Size and type	35
4.4.2 Placement	38
4.5 Accuracy analysis	39
4.6 Photogrammetry algorithm	41
5 UWB setup	47
5.1 Basic principle	47

5.2	Pozyx sensors	47
5.3	Calibration	49
5.4	Range test	50
6	Wind speed & angular measurements	52
6.1	Apparent wind speed measurements	52
6.1.1	Experimental setup	52
6.1.2	Calibration and directional sensitivity	53
6.2	Angular measurements	55
6.2.1	Calibration	56
7	Experimental design	57
7.1	Experimental setup	57
7.2	KCU	59
7.2.1	Definition of inputs	60
7.3	Intrusiveness	60
7.3.1	Flow obstruction	61
7.3.2	Mass	61
7.4	Time-synchronisation	61
7.4.1	Camera synchronisation	62
7.4.2	KCU and UWB synchronisation	64
7.4.3	Photogrammetry and UWB synchronisation	64
7.5	Measurement architecture	66
8	Post-processing	69
8.1	General post-processing operations	69
8.1.1	(I) Linear calibration of raw signals	69
8.1.2	(II) Handling missing samples	69
8.1.3	(III) Reducing high-frequency noise	70
8.1.4	Overview of applied post-processing per signal	71
8.2	Photogrammetry-specific post-processing	72
8.2.1	(i) Correction for camera-rig bending and twist	73
8.2.2	(ii) Transformation to a wing-fixed reference frame	81
8.2.3	(iii) Outlier rejection in the 3D marker cloud	82
8.2.4	(iv) Wing shape reconstruction	83
9	Results	85
9.1	In-flight vs. CAD geometry	85
9.2	Billowing and twist	86
9.2.1	Billowing	87
9.2.2	Twist angle	88
9.3	(De)power deformation	89
9.4	Steering-induced deformations	93
10	Structural model validation	98
10.1	Structural finite element model	98
10.2	The hanging kite	99
10.3	FEM validation criterion	101
10.4	Results	101
11	Conclusions	103
12	Recommendations	105
12.1	Improvements to the current methodology	105
12.2	Future research directions	106

Appendix A List of requirements	107
A.1 Experimental requirements	107
A.2 Instrumental requirements	107
A.3 Data processing and validation requirements	109
A.4 Logistical and practical requirements	110
Appendix B Sensor trade-off	111
B.1 Global shape mapping technologies	111
B.2 Time-of-Flight (ToF) and active ranging technologies	113
B.3 Structural and kinematic sensing technologies	113
Appendix C Fresnel zone calculation	114
Appendix D Improved vane design	115
Appendix E Custom printed parts	116
Appendix F Bridle geometry flight tests	118
Appendix G Bridle geometry static test	121
Appendix H Additional results	123
H.1 Left turn	123
H.2 Automated video results	124
H.3 Static test results	124
References	126

List of Figures

1	Me flying a kite on the beach (Texel, The Netherlands, 2010).	i
1.1	Schematic representation of the pumping cycle in crosswind flight [9]. During reel-out, the kite is powered and flies at high speed in specific patterns, generating high tether tension that drives the ground-based generator. At maximum extension, the kite is depowered and reeled in under low tension, after which the cycle restarts. The net effect is positive energy generation.	2
1.2	Soft-wing leading-edge inflatable (LEI) ground-based AWES [10], based on the TU Delft V3 kite, further developed by Kitepower B.V. [11]. (a) Overview of the main system components: wing, bridle system, KCU, tether, and ground station. (b) Front view of the kite and bridle system, highlighting control inputs. The rear bridles (red) are adjusted via the KCU: symmetric actuation of the depower tape modifies the effective angle of attack, while asymmetric steering inputs induce turning manoeuvres.	2
1.3	TU Delft V3 kite used in this study, with different bridle layout. (a) Two-view illustration of the baseline V3. (b) Experimental configuration without an airborne KCU [10].	3
2.1	Flow topology around a LEI kite, showing pressure-side recirculation due to the inflated leading edge and a suction-side laminar separation bubble [21].	7
2.2	Visual evidence of three-dimensional flow structures on kites using smoke visualization techniques. (a) The tips show coherent tip vortices on a RAM-air kite [27]. (b) Pressure-side separation is visible on a LEI kite [17].	7
2.3	Force equilibrium in static flight [29]. Here F_T is the tether force, mg the weight of the kite, L and D the lift and drag, \mathbf{v}_w the wind speed vector and β the kite elevation angle.	8
2.4	Coordinate systems for a tethered kite: ground frame, spherical position (β, ϕ, r) , local tether frame (x_t, y_t, z_t) , and heading ψ [16].	9
2.5	Coupled aero-structural model of the TU Delft V3 kite [10]. Left: distributed mass-spring-damper network for the structure, with nodes placed at bridle attachment and structural junctions. Right: vortex step method on a segmented lifting surface conforming to the deformed shape.	9
2.6	Depowering: shape of powered (red) and depowered (blue) kite [8].	11
2.7	Steering: shape before (left) and after (right) a steering input [17].	11
2.8	Common bar-based definitions of (a) depower and (b) steering input [38]. In a KCU, the same inputs are implemented as front-back and left-right line-length changes. Note that steering input in the image is defined as δ , whilst u_s is used throughout this report. Additionally, the power setting u_p is different from depower setting u_{dp} used in this research.	12
2.9	Global and local vibration modes of LEI kites and their characteristic time scales. (a) Mode taxonomy highlighting billowing, jellyfishing, bunny-ear flapping and leading-edge indentation [39]. (b) Comparison of deformation and flight-path frequencies indicating that most global deformation energy lies below about 2 Hz [31].	12
2.10	Single-camera photogrammetry setup used by Hanke & Schenk in flight. The setup is one of the few visual tracking setups used whilst flying.	14
2.11	Wireframe model as described by Jonard [15]	16
2.12	Bar construction designed by Oehler [49] to carry the Pitot tube and other measurement units, and perform wind speed measurements on the flying kite.	17
2.13	Most recent research into kite deformations using stereoscopic photogrammetry: the experimental ground stereoscopic photogrammetry setup used in [15].	18
4.1	Single-camera photogrammetry principle. A point can be tracked in time, but only in 2D. Depth cannot be resolved, as the 2D image plane provides only the x and y coordinates. Determining z requires an additional view of the object.	24

4.2	Triangulation principle in stereoscopic photogrammetry. After image rectification, discussed later in this chapter, corresponding points in both camera views lie on the same epipolar line. This geometric constraint reduces the problem to two dimensions and enables the 3D coordinates to be solved through simple geometrics.	25
4.3	GoPro Hero9 Black from three different angles. Two such cameras are used in the stereoscopic photogrammetry setup [116].	26
4.4	Intrinsic calibration parameters estimated in OpenCV: focal lengths (f_x, f_y) , principal point (c_x, c_y) , and distortion coefficients describing deviations from the ideal pinhole projection.	27
4.5	Calibration patterns used in this study [118]. These patterns were printed on A3 paper to maximise visibility.	29
4.6	Effect of lens distortion correction on a chequerboard calibration image. (a) Corner detection in OpenCV on the chequerboard pattern (GoPro Hero9 Black, 5K wide-angle). Barrel distortion is visible near the image edges. (b) Undistorted chequerboard using the estimated intrinsics and distortion coefficients (same frame as (a)). Barrel distortion is removed; the result is cropped to a rectangular aspect.	30
4.7	Baseline B and disparity concept. For a fixed setup, the closer point at distance z_1 yields a larger horizontal disparity $\delta_{x,1}$ than the further point at z_2 with horizontal disparity $\delta_{x,2}$. Increasing the baseline magnifies horizontal disparities for all depths, thereby improving depth resolution.	32
4.8	Estimated depth error ϵ_z at $z = 7$ m vs. baseline B for several disparity errors $\epsilon_{\delta,x}$. The marked point indicates the design minimum baseline that just meets REQ-EXP-01b in the worst case.	33
4.9	Example of the stereoscopic calibration input images (chequerboard). This is one of the pairs used in OpenCV's <code>stereoCalibrate</code> to estimate the relative rotation \mathbf{R} and translation \mathbf{t} between the cameras.	34
4.10	Rectification pipeline, adapted from [125]. (a) Raw stereo images; (b) undistortion removes lens distortion; (c) rectification aligns epipolar lines horizontally; (d) the images are cropped to retain rectangular frames.	34
4.11	Markers tested for detection analysis. The minimum pixel size for detection of each of these markers was determined through experiments. This pixel size could then be translated to a required marker size at distance z	36
4.12	Pixel scaling at a given distance z . α_H and α_V denote the horizontal and vertical FOV of the wide lens. The physical size represented by one pixel is px_x and px_y . For an ArUco, s_m is the side length; for a circle, s_m is the diameter. To be conservative, the conversion to millimetres uses $\max(px_x, px_y)$	36
4.13	Principle of sub-pixel accuracy; circular markers used as example. Boundary pixels contain a blend of marker and background colour, allowing the true marker centre to be estimated more precisely than one pixel.	37
4.14	Red cross markers on both white (left) and black (right) background.	38
4.15	Markers used and their placement on the V3. Markers on the struts have a black background, whereas markers on the LE have a white background.	38
4.16	Schematic of the accuracy test. Markers (red crosses) are placed on a flat wall at a distance $\Delta z \approx 6$ m. Markers are divided by even spacings Δx and Δy . The geometry reflects the characteristic camera-marker distances in the experiments.	39
4.17	Coupled, rectified frames of the accuracy test. Identified markers are shown and numbered; the trapezium shape is illustrated for clarity.	40
4.18	Example grid of reconstructed markers (blue) and their corresponding points on the fitted trapezium of size $3\Delta x \times \Delta y$ (red). Values shown are Euclidean errors ϵ . Note that the z axis shows centimetres.	41
4.19	Side-by-side comparison of original image (left) and image with adjusted brightness and contrast (right). With increased brightness and contrast, features on the kite become much more visible, whilst noise is removed as the background looks almost transparent.	42

4.20	Illustration of epipolar constraints and cost matrix construction. For a detected marker in the left frame (\mathbf{x}_i), a bounding box is defined by the vertical tolerance $\delta_{y,\max}$ and horizontal disparity limits $\delta_{x,\min}$ and $\delta_{x,\max}$. Candidate matches in the right frame are only considered if they fall inside this bounding box. A marker within the box (e.g. \mathbf{x}'_j) is assigned a finite cost in the matrix equal to its Euclidean distance from \mathbf{x}_i , whereas a marker outside the box (e.g. \mathbf{x}'_k) is penalised with an infinite cost. This procedure ensures that only geometrically consistent correspondences are retained for stereo matching.	44
4.21	Matched markers and their epipolar lines in both the left and right frame after filtering.	45
5.1	UWB sensors placed on the tips of the wing, measuring the span b	47
5.2	Pozyx sensors used for the experiment: (a) Pozyx shield v1.5, used for sending commands and subsequently receiving ranging data. (b) Pozyx anchor v1.4, used for ranging between one another. The relevant components of both sensors are highlighted.	48
5.3	The error (ϵ_{UWB}) of the UWB sensors before and after calibration, for distances ranging from 5 to 15 meters.	49
5.4	Range test between a Pozyx anchor (v1.4) and Pozyx shield (v1.5). The communication link was lost at ≈ 25 m, insufficient for ground-based operation.	50
5.5	Raspberry Pi 3A+ used to run the UWB ranging code and store data during flight. Only relevant subsystems have been highlighted.	51
6.1	Pitot tube principle. The free stream velocity V_∞ flows in the front, leading to a total pressure p_{tot} measured there. Perpendicular to the flow, the flow velocity is equal to $V = 0 \text{ ms}^{-1}$, with static pressure p_s	52
6.2	Wind Tunnel test setup, with measurements conducted (a) inside and (b) outside the wind tunnel. Here, V_a is the magnitude of the measured apparent wind speed.	53
6.3	The experimental setup inside the wind tunnel. The casing (in white) has two vanes for angular measurements and the Pitot tube on the left.	53
6.4	Protractor and 3D-printed indicator used to accurately determine the inflow angle ϕ , mounted on the rotating bar.	53
6.5	Measured wind speed for different inflow angles ϕ before and after calibration.	54
6.6	Working principle of the angular vanes attached to the pitot casing.	55
6.7	Angular vane measurements before and after calibration. The calibration clearly improves the results of the angular measurements, which overshoots for larger angles of attack, particularly at higher wind speeds.	56
7.1	Schematic of the experimental setup. A rigid Γ -bar, mounted in the bridles beneath the wing, carries two cameras separated by a stereo baseline B , a Pitot tube, and a sensor box (including the UWB receiver). UWB ranging sensors are installed at both wingtips to infer the instantaneous span b . The cameras' fields of view are directed toward the lower surface of the wing.	57
7.2	Experimental setup: (a) Γ -bar detail; (b) overall in-flight configuration.	58
7.3	Measured α_{fl} (front-line inclination) versus aerodynamic angle at the wing α_w	59
7.4	Front view of KCU actuation; depower and steering tapes through motors.	59
7.5	Audio-based synchronisation of GoPro recordings. (a) Original, unsynchronised signals. (b) After lag correction.	63
7.6	Synchronised sound waves from the final experiment in Bangor-Erris.	63
7.7	Measured drift between two GoPro cameras over a 15-minute period. Values remain far below the frame-mismatch threshold.	64
7.8	Flash sequence used to map GoPro frames to UTC.	64
7.9	Flash onset detection via brightness slope. The vertical line marks the synchronisation frame.	65
7.10	Lag between LED flash (true UTC) and the first UWB UTC log. Mean offset is 1.736 s.	66

7.11 Schematic overview of the experimental setup. The figure illustrates the full measurement architecture, including sensors, data flows, and synchronisation links. The GS issues commands to the KCU, which actuates the kite and interfaces with the Pitot tube and angular vanes. Wing shape is measured independently by the photogrammetry subsystem and the UWB subsystem. All outputs are postprocessed and synchronised to UTC to form the final dataset.	67
8.1 Example of zero-phase EMA filtering on UWB span data. The mean distance has been subtracted to show only fluctuations around the mean. The raw measurements (blue) exhibit high-frequency jitter, while the zero-phase EMA output (orange, $\alpha = 0.9$) provides a smoother trace without apparent time shift.	72
8.2 Effect of yaw change due to bending of the Γ -bar. (a) Nominal stereoscopic calibration ($\Delta\psi_{\text{bend}} = \Delta\psi_{\text{corr}} = 0$); (b) converging cameras ($\Delta\psi_{\text{bend}} < 0$, $\Delta\psi_{\text{corr}} > 0$) resulting in overestimated depth and enlarged scale; and (c) diverging cameras ($\Delta\psi_{\text{bend}} > 0$, $\Delta\psi_{\text{corr}} < 0$) resulting in underestimated depth and reduced scale. The correction angles $\Delta\psi_{\text{corr}}$ are applied in post-processing with opposite sign to the deformation-induced misalignment.	74
8.3 Effect of roll change caused by asymmetric bridle loading. (a) Nominal calibration ($\Delta\phi_{\text{bend}} = \Delta\phi_{\text{corr}} = 0$); (b) clockwise twist of the bar ($\Delta\phi_{\text{bend}} < 0$, $\Delta\phi_{\text{corr}} > 0$) producing opposite vertical disparities on either side of the frame; and (c) counter-clockwise twist ($\Delta\phi_{\text{bend}} > 0$, $\Delta\phi_{\text{corr}} < 0$) producing the inverse effect.	75
8.4 Effect of pitch change due to torsional deformation of the Γ -bar. (a) Nominal calibration ($\Delta\theta_{\text{bend}} = \Delta\theta_{\text{corr}} = 0$); (b) upward pitch ($\Delta\theta_{\text{bend}} < 0$, $\Delta\theta_{\text{corr}} > 0$) causing uniform downward marker shift and positive vertical disparity; and (c) downward pitch ($\Delta\theta_{\text{bend}} > 0$, $\Delta\theta_{\text{corr}} < 0$) causing uniform upward shift and negative vertical disparity.	75
8.5 Three-dimensional parameter study showing the influence of small angular corrections on reconstruction quality. The colour scale indicates (left) the average vertical disparity $ \delta_y $ and (right) the span deviation Δb between photogrammetry and UWB measurements. The results confirm that span errors are predominantly driven by yaw correction ($\Delta\psi_{\text{corr}}$), while roll ($\Delta\phi_{\text{corr}}$) and pitch ($\Delta\theta_{\text{corr}}$) primarily influence vertical alignment.	78
8.6 Parameter slice through the optimized yaw correction $\Delta\psi_{\text{corr}}^*$. (Left) Mean vertical disparity $ \delta_y $ as function of roll and pitch corrections; (Right) corresponding span deviation Δb compared to the UWB reference. Roll and pitch corrections mainly affect vertical alignment, while span variations are dominated by yaw.	79
8.7 Camera and wing reference frames. The dense sets of markers on the two centre struts (strut 3 and strut 4) define a wing-fixed frame at each instant: the origin \mathbf{o}_w is located midway between the strut centroids, the x -axis $\hat{\mathbf{x}}_w$ follows the average direction of the two centre struts, the y -axis $\hat{\mathbf{y}}_w$ is spanwise, and the z -axis $\hat{\mathbf{z}}_w$ is normal to the wing plane.	82
8.8 Marker indexing on the wing during flight. The distinction between the leading-edge (LE) markers and the strut markers is indicated, with the strut markers numbered from 0 to 7. Some markers remain unidentified, either because they do not pass the detection filters or because they are occluded in either of the camera views.	84
9.1 Comparison between the powered in-flight kite geometry (blue) and the undeformed CAD representation (grey).	86
9.2 Segments considered for billowing analysis (in white). This video still is taken during a right turn.	87
9.3 Length of kite segments, indicating billowing for all flight cases.	87
9.4 Definition of (a) twist angle τ and (b) anhedral angle γ	88
9.5 Spanwise twist angle τ for several flight states.	89
9.6 Characteristic geometry change between a powered and a depowered kite.	90
9.7 Side view of powered/depowered state in camera coordinate frame. The depowered state (orange) clearly exhibits a pitch forward with respect to the powered (blue) state.	91

9.8	Shape change for the claims made, all in the wing-fixed reference frame. The identifiers correspond to the claims: (i) and (iii) result in a yaw motion (bottom view), and (ii) causes a roll motion (front view). The variable local angle in (i) is a consequence of the variable twist (see Fig. 9.5).	94
9.9	Shape change during a right-turn manoeuvre in the wing reference frame.	95
9.10	Bottom and front view of the right-turn geometry visualized in the camera reference frame. 96	
10.1	Kite composition in the FEM [142].	98
10.2	Hanging-kite experimental setup used for static deformation measurements.	99
10.3	Overview of the load cases applied in the static hanging-kite experiment.	100
C.1	First Fresnel zone geometry. The maximum radius at the midpoint defines the required ground clearance for maintaining unobstructed line-of-sight communication.	114
D.1	Improved angular vane design (not to scale). (a) original flat vane located close to the casing; (b) redesigned vane with symmetric airfoil section and increased stand-off from the casing.	115
E.1	Custom 3D-printed components used in the experimental setup. All parts were manufactured in PETG HF with 35% gyroid infill to balance stiffness, strength, and weight. . . .	117
F.1	Bridle attachment points on the wing. LE connections (left) correspond to the bridle system shown in Fig. F.2; TE connections (right) correspond to Fig. F.3.	118
F.2	Schematic of the leading-edge bridle system and corresponding segment numbering. . . .	119
F.3	Schematic of the trailing-edge bridle system and corresponding segment numbering. . . .	120
G.1	Leading-edge bridle layout used during the static (upside-down) tests.	122
G.2	Trailing-edge bridle layout used during the static (upside-down) tests.	122
H.1	Shape deformation during a left-hand turning manoeuvre.	123
H.2	Example snapshot from the automated result video, showing four point cloud views, synchronised video footage, and flight-state overlays.	124
H.3	Static test results for pressure level P1.	125
H.4	Static test results for pressure level P2.	125

List of Tables

1	List of Roman symbols used throughout the report.	xiii
2	List of Greek symbols used throughout the report.	xiv
3	List of subscripts used throughout the report.	xv
4	List of superscripts used throughout the report.	xvi
5	List of abbreviations used throughout the report.	xvii
1.1	Design geometry of the V3 kite.	3
4.1	Relevant parameters of the GoPro Hero9 Black. Only specifications relevant to this study are listed, considering the highest video quality setting to ensure maximum accuracy [117].	26
4.2	Specifications of the calibration patterns used.	29
4.3	Calibration results for camera 1. The reprojection error confirms that calibration with the chequerboard pattern was the most accurate.	31
4.4	Calibration results for camera 2. The reprojection error again shows that the chequerboard pattern produced the most reliable results.	31
4.5	Final intrinsic parameters for both cameras	31
4.6	Baseline for the V3 test cases at $z = 7$ m. A value of $SF = 2$ was applied to the computed minimum.	33
4.7	Stereo extrinsics for the adopted setup ($B = 84$ cm), rounded to three decimals.	35
4.8	Stereo calibration quality metric (averaged across calibration images).	35
4.9	Minimum detectable sizes for each marker type. Millimetre sizes are computed at the design distance $z = 7$ m using Fig. 4.12.	37
4.10	Marker counts per strut on the V3 (measured from the centreline outward).	39
4.11	Metrics used in the accuracy tests described in Fig. 4.16.	40
4.12	Errors for the stereoscopic setup at $\Delta d_{\max} \approx 7$ m.	41
4.13	Constraining limits used in marker matching. These values create a bounding box based on a marker in the left frame, stating that the corresponding marker in the right frame should be positioned in this box. The vertical tolerance $ \delta_{y,\max} $ is set rather high as a result of bar bending, discussed in Section 8.2.1.	43
4.14	Inner and outer radius used for background colour definition.	45
5.1	Relevant parameters of the Pozyx sensors. Only specifications relevant to this study are listed. Values on accuracy and range seem to differ greatly, and are therefore analysed independently for the specific set of sensors and setup.	48
5.2	Calibration values for the UWB sensors	49
6.1	Calibration constants used for the Pitot tube.	54
6.2	Average absolute and relative error for the Pitot tube.	55
6.3	Calibration constants used for the angular vanes.	56
7.1	Depower settings.	60
7.2	Steering settings.	60
7.3	Mass contribution of each item of the experimental setup. Only airborne items are listed; total masses below 15 g are neglected.	62
7.4	Model outputs, all subsystems synchronised and combined	68
8.1	Overview of post-processing operations applied to each data stream.	72
8.2	Summary of the qualitative influence of small angular misalignments on disparity and reconstructed geometry.	76
8.3	Reference distances between outermost markers on the carbon-reinforced struts used for adjusted accuracy estimation.	80

9.1	Measured reduction in span between the undeformed CAD geometry and the powered in-flight geometry.	86
9.2	Difference in strut-to-strut distance with respect to the CAD model for several flight states.	88
9.3	In- and outputs corresponding to the two flight states provided in Fig. 9.6	90
9.4	Summary of relative findings when going from fully powered to fully depowered state. Uncertainties are based on the average photogrammetry and UWB accuracies, found in Chapter 4 and Chapter 5 respectively	93
9.5	In- and outputs for straight flight and the right turn shown in Fig. 9.9.	95
9.6	Summary of findings when going from straight flight to right turn (powered). Uncertainties are based on the average photogrammetry and UWB accuracies, found in Chapter 4 and Chapter 5 respectively.	97
10.1	Combined overview of static test cases with angle adjustments and applied loads.	100
10.2	Shape correlation coefficients and mean absolute distance errors for all static load cases.	102
B.1	Trade-off of different measurement techniques for mapping the entire kite, measuring node position/distance or structural/kinematic sensing technologies. The trade-off criteria are derived from the requirements. * Accuracy is judged for the stated purpose and range. For IMUs, strain tapes and flex sensors, positional accuracy is not considered. ** Despite possibly being expensive, these units are readily available for use.	112
B.2	Legend for Trade-Off table	112
F.1	Bridle length table used during the flight experiments.	118
G.1	Bridle length table used during the static experiments.	121
H.1	Inputs and outputs corresponding to the two flight states shown in Fig. H.1.	123

List of Symbols and Abbreviations

Roman Symbols

Table 1: List of Roman symbols used throughout the report.

Symbol	Definition	Unit
\mathbf{A}	Assignment matrix	–
A	Projected wing area	m^2
\mathcal{A}_i	Annulus region around marker i	px
a, b	Calibration constants	–
B	Stereo camera baseline	m
B_i	Image brightness of pixel i	–
b	Wingspan	m
\mathbf{C}	Cost matrix	–
C	Cluster	–
$C_p(x)$	Pressure coefficient distribution along the chord	–
C_R	Resultant aerodynamic force coefficient	–
c	Chord / speed of light	$\text{m} / \text{m s}^{-1}$
ce	Centroid	px
c_x, c_y	Optical centre	px
D	Aerodynamic drag	N
d	Distance	px / m
\mathbf{d}	Distortion vector	–
d_1, d_2	Geometric distances in calibration / range relations	m
drift	Time drift	s
$e^{j\omega}$	Complex frequency variable on the unit circle in the z -domain	–
\mathbf{F}	Fundamental matrix	–
F_T	Tether force	N
f	Frequency	Hz
fps	Frames per second	s^{-1}
f_x, f_y	Focal point	px
g	Gravitational acceleration	m s^{-2}
H	Image height	px
$H(e^{j\omega})$	Frequency response of the forward EMA filter	–
$H_0(e^{j\omega})$	Zero-phase EMA filter transfer function	–
h	Wing height	m
\mathbf{I}	Identity matrix	–
I	Image intensity	–
j^*	Index of closest predecessor marker in nearest-neighbour matching	–
k_1, k_2, k_3	Radial distortion coefficients	–
k_s	Steering gain	mm
\mathbf{K}	Camera intrinsic matrix	–
lag	Time lag	s
L	Aerodynamic lift / Length	N / m
L_n	Cost term in reconstruction / optimisation	–
$\mathcal{L}(\pi)$	LE in wireframe model	–
$\mathcal{L}(t)$	Strut in wireframe model	–
m	Mass (generic) / finite number	$\text{kg} / \text{–}$

Symbol	Definition	Unit
M_{ij}	Element (i, j) of a matrix (e.g. cost matrix) / Spatial image moment of order (i, j) for a contour	- / px ^{$i+j$}
N	Finite number	-
n	Finite number	-
\mathbf{o}	Origin	-
p	Pressure	Pa
p_1, p_2	Tangential distortion coefficients	-
$\mathbf{p}(t)$	3D point / marker position as function of time	m
\mathbf{p}	3D point	m
\mathbf{R}	Rotation matrix	-
\mathbb{R}	Set of real numbers	-
\mathbb{R}^n	nD real vector space (Euclidean plane)	-
\mathbf{Q}	Centred point matrix	-
r	Radial distance (e.g. spherical coordinate)	m
Re	Reynolds number	-
\mathbf{S}	Singular value matrix of \mathbf{Q}	-
S	Flattened surface area	m ²
s	Pixel spacing / size	px / m
SC	Shape correlation factor	-
SF	Safety factor	-
T	Period	s
t	Time / thickness	s / m
\mathbf{t}	Translation matrix	-
th_{tape}	Tape thickness	mm
\mathbf{U}	Left singular vectors of \mathbf{Q}	-
u_{dp}	Depower input (KCU command)	-
u_s	Steering input (KCU command)	-
\mathbf{V}	Right singular vectors of \mathbf{Q}	-
V	Velocity (generic)	m s ⁻¹
$V(x, y)$	Value (brightness) channel of the HSV image at pixel (x, y)	-
V_a	Apparent wind speed at the kite	m s ⁻¹
V_w	Ambient wind speed	m s ⁻¹
$\hat{\mathbf{v}}$	First right singular vector defining the best-fit line direction	-
w	Projected wingspan	m
W	Image width	px
x, y, z	Cartesian coordinates in 3D space	m
x, y	Camera frame coordinates	px
X, Y, Z	Coordinates in world frame	m
\mathbb{Z}^2	2-D integer pixel coordinate space	-
z	Depth	m

Greek Symbols / Others

Table 2: List of Greek symbols used throughout the report.

Symbol	Definition	Unit
α	Angle of attack / smoothing factor in EMA filtering / gain	deg / - / -
β	Elevation angle (spherical coordinates) / bias	deg / -
β_s	Sideslip angle	deg
Γ	Shape of bar	-
γ	Anhedral angle	deg

Symbol	Definition	Unit
Δ	Finite change / difference operator	–
δ	Disparity	px
ϵ	Error term (e.g. reconstruction or measurement error)	–
ζ	Bottom shear angle	deg
θ	Pitch angle	deg
κ	Slope constant	–
λ	Wavelength	μm
μ	Correction angle	deg
ν	Median background brightness	–
ξ	Front shear angle	deg
ϕ	Roll angle / azimuth / inflow angle	deg
Φ	Shape vector	–
ψ	Yaw angle / heading angle	deg
ρ	Air density	kg m^{-3}
σ	Standard deviation	–
τ	Twist angle	deg
ω	Angular frequency	rad s^{-1}
\emptyset	Diameter	m
$\ \cdot\ _2$	Euclidean (L2) norm	m

Subscripts

Table 3: List of subscripts used throughout the report.

Subscript	Meaning
a	Aerodynamic / apparent (e.g. V_a)
actual	Actual quantity
adj	Adjusted quantity (e.g. adjusted accuracy)
anchor	Anchor-specific quantity
av	Angular vane related
avg	Averaged quantity
bend	Bending quantity
c	Camera frame
cam	Camera-specific quantity
CB	Billowing
calibrated	Calibrated value
corr	Corrected quantity
CS-T	Centre-strut to tip
d	Depth-related
depowered	Depowered configuration
dis	Distorted (before lens distortion correction)
dynamic	Dynamic (e.g. p_{dynamic})
e	Epipolar-related
exp	Experimental value
fl	Front lines (e.g. α_{fl})
frame	Index of video frame
H	Horizontal (e.g. field of view)
i	Index i
image	Image-plane quantity
in	Inner
ip	In-plane component

Subscript	Meaning
k	Index k
KCU	Belonging to the Kite Control Unit
L	Left (camera / line)
left	Left (camera / line)
lim	Limit value (e.g. r_{lim})
local	Local value
m	Marker-related
match	Used in matching condition (tracking)
max	Maximum value
measured	Measured quantity
min	Minimum value
mod	Model
n	Index or number n
op	Out-of-plane component
out	Outer
p	Photogrammetry / plane frame
PL	Point load related
powered	Powered configuration
R	Right (camera / line)
right	Right (camera / line)
r	Reprojection
raw	Raw value (before calibration)
s	Static / steering / sampling / stagnation
start	Start (e.g. initial value)
T	Tether-related quantity
t	Time-related / tether frame (literature only)
tan	Tangential
TL	Tip load related
tot	Total quantity
thr	Threshold
u	Undistorted (after lens correction)
UWB	Ultra-wideband-related quantity
V	Vertical (e.g. field of view)
w	Wing frame / wind (literature only)
wc	Wing-to-camera transform (e.g. \mathbf{R}_{wc})
ws	Wind-speed related
x, y, z	3D coordinate direction
∞	Free-stream (e.g. V_{∞})
δ	Disparity-related
π	Leading edge related

Superscripts

Table 4: List of superscripts used throughout the report.

Superscript	Meaning
b	Backward filtering pass (EMA)
f	Forward filtering pass (EMA)
*	Optimal / selected element (e.g. best match)
'	Updated / intermediate value
T	Matrix transpose

Abbreviations

Table 5: List of abbreviations used throughout the report.

Abbrev.	Meaning
1D	One-dimensional
2D	Two-dimensional
3D	Three-dimensional
3D-GS	3D Gaussian splatting
AoA	Angle of attack
AoS	Angle of sideslip
AWE	Airborne wind energy
AWES	Airborne wind energy system
B.V.	Besloten vennootschap (Dutch private limited company)
CAD	Computer-aided design
CDGPS	Carrier-differential GPS
CFD	Computational fluid dynamics
DIC	Digital image correlation
DPV	Data processing and validation
EKF	Extended Kalman filter
EMA	Exponential moving average
EMI	Electromagnetic interference
EXP	Experimental
FEM	Finite element method
FBG	Fibre Bragg grating
FOSS	Fibre optic shape sensors
FOV	Field of view
FPS	Frames per second
FSI	Fluid-structure interaction
GNSS	Global navigation satellite system
GPS	Global positioning system
GS	Ground station
HWA	Hot-wire anemometry
IIR	Infinite impulse response
IMU	Inertial measurement unit
INS	Instrumental
KCU	Kite control unit
KLT	Kanade-Lucas-Tomasi tracker
LED	Light-emitting diode
LE	Leading edge
LEI	Leading-edge inflatable
LiDAR	Light detection and ranging
LP	Logistical and practical
LS-DYNA	Livermore software dynamic nonlinear analysis
ND	Not determined
NRTK	Network real-time kinematic
NTP	Network time protocol
PC	Personal computer
PETG HF	Polyethylene terephthalate glycol, high flow
PIV	Particle image velocimetry
PSM	Particle system model
RANS	Reynolds-Averaged Navier-Stokes
REQ	Requirement

Abbrev.	Meaning
RQ	Research question
ROI	Region of interest
RTK GPS	Real-time kinematic GPS
SALMA	Single-anchor localisation system
SSE	Sum of squared errors
SSH	Secure shell
STVD	Standard deviation
SVD	Singular value decomposition
TAS	True airspeed
TE	Trailing edge
ToF	Time-of-flight
UWB	Ultra-wideband
UTC	Coordinated universal time
VSM	Vortex step method
VWT	Virtual wind tunnel

Introduction 1

Global temperatures are rising, and even modest increases can trigger profound and often irreversible changes in the Earth's systems. Scientific research has shown that warming of approximately 1.5 °C above pre-industrial levels is sufficient to significantly increase the frequency and severity of heatwaves, intensify tropical storms, accelerate sea-level rise, and place considerable stress on ecosystems and biodiversity [1]. Such changes also have far-reaching consequences for human societies, including threats to food production and water security in regions already vulnerable to scarcity.

The primary driver of this warming is human activity, particularly the large-scale combustion of fossil fuels such as coal, oil, and natural gas. These energy sources are responsible for the majority of greenhouse gas emissions, which accumulate in the atmosphere and disrupt the balance of the global climate [2]. Unless emissions are rapidly reduced, the risks of runaway warming, ice-sheet collapse, and widespread biodiversity loss will only intensify, leading to detrimental consequences.

Addressing these challenges requires a fundamental transformation of how energy is produced and consumed. Expanding and accelerating the adoption of renewable energy technologies is one of the most effective strategies to reduce emissions and mitigate climate impacts. Recent research highlights that countries that have scaled up renewable energy use have already seen measurable declines in carbon dioxide emissions, underscoring the urgent need to continue developing and deploying innovative renewable solutions worldwide [3].

1.1 Airborne wind energy & the pumping cycle

Airborne wind energy (AWE) has emerged as a promising alternative to conventional wind turbines, primarily because it allows for harvesting wind at higher altitudes, where conditions are stronger and more consistent. By employing tethered flying devices instead of rigid towers, AWE systems can operate beyond the reach of traditional turbines. Furthermore, they require up to 90% less construction material [4], reducing both their environmental footprint and infrastructure costs [5]. These characteristics make AWE particularly attractive for decentralized energy production and for use in remote locations without access to the electrical grid [6].

The concept of using kites for large-scale power generation was famously analysed by Miles L. Loyd in 1980 [7]. Loyd demonstrated that by flying kites crosswind — at high velocity transverse to the wind direction — substantial lift forces are generated and converted into usable power, the reason crosswind flight is so effective. His analysis showed that such a system could outperform conventional wind turbines in power output per unit of material, due to both higher wind speeds at altitude and the absence of a tower.

Among the various AWE concepts, the so-called pumping cycle has shown the most promise for practical implementation. In this approach, a tethered kite is flown in crosswind patterns, typically figure-of-eight manoeuvres, to maximise aerodynamic forces on the tether. The pumping cycle consists of two alternating phases. During the reel-out phase, the kite is powered, flying crosswind at high speed. The resulting lift creates a strong tether force, which drives a ground-based drum-generator system and produces electricity. Once the tether has reached its maximum extension, the system transitions into the reel-in phase, where the kite is depowered by adjusting its orientation with respect to the flow. This reduces the aerodynamic load, allowing the tether to be reeled in with minimal energy consumption [8, 9]. Since the energy extracted in the reel-out phase exceeds the energy required for retraction, the cycle yields a net positive energy output (see Fig. 1.1).

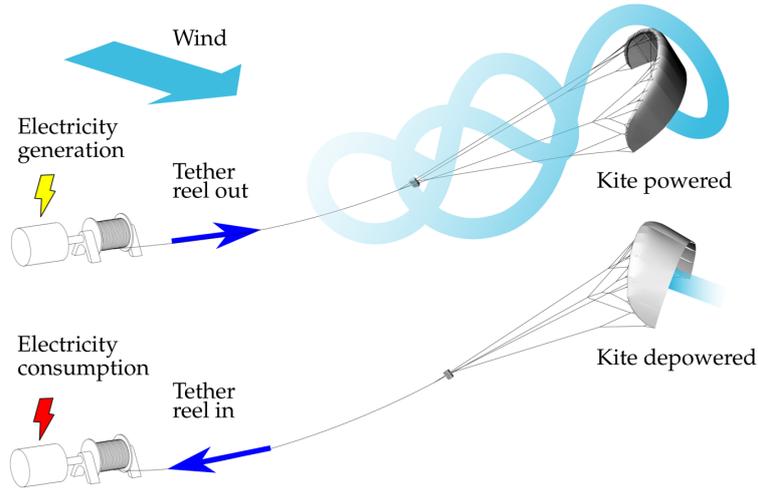


Figure 1.1: Schematic representation of the pumping cycle in crosswind flight [9]. During reel-out, the kite is powered and flies at high speed in specific patterns, generating high tether tension that drives the ground-based generator. At maximum extension, the kite is depowered and reeled in under low tension, after which the cycle restarts. The net effect is positive energy generation.

A soft-wing single-tether pumping cycle system is composed of five key components: (i) the wing, which provides aerodynamic lift; (ii) the bridle system, which distributes aerodynamic loads and enables steering and depowering; (iii) the kite control unit (KCU), an automated device suspended beneath the wing that regulates depower and steering inputs; (iv) the tether, which transmits mechanical power to the ground; and (v) the ground station, typically a drum-generator module that converts mechanical work into electricity [10]. These components are illustrated in Fig. 1.2.

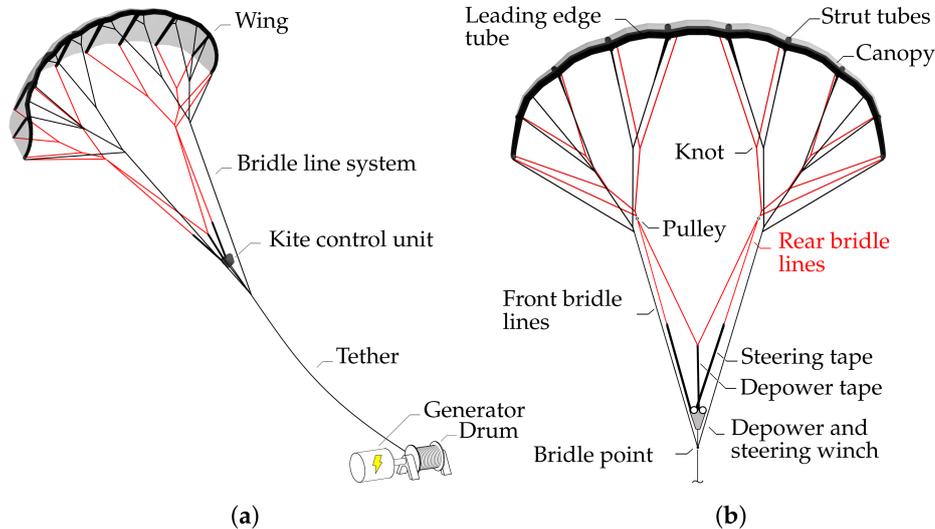


Figure 1.2: Soft-wing leading-edge inflatable (LEI) ground-based AWES [10], based on the TU Delft V3 kite, further developed by Kitepower B.V. [11]. (a) Overview of the main system components: wing, bridle system, KCU, tether, and ground station. (b) Front view of the kite and bridle system, highlighting control inputs. The rear bridles (red) are adjusted via the KCU: symmetric actuation of the depower tape modifies the effective angle of attack, while asymmetric steering inputs induce turning manoeuvres.

Through the interaction of these components, continuous automated pumping cycles can be executed, enabling efficient and scalable power generation. While current developments often focus on soft leading-edge inflatable (LEI) kites, other variations of airborne wind energy systems have been investigated, such as drag-based designs, airborne turbines with on-board generation, and moving ground stations. Despite this diversity, the pumping cycle remains the most mature and widely studied approach for large-scale deployment [5].

1.2 The V3 kite

The system considered in this work employs a soft LEI wing, attached via a bridle system and lines to the KCU for control purposes. This entire flying system is connected to the ground station through a tether. The kite used is the TU Delft V3 wing, further developed by Kitepower B.V. [11]. The V3 kite is selected as the research platform due to the availability of an established dataset, enabling meaningful cross-comparison between new experiments and existing numerical and experimental results. Expanding this dataset with the present tests strengthens the evidence base and facilitates validation of modelling approaches. Furthermore, comparable LEI kites are currently in commercial operation. A graphical depiction of the V3 kite is shown in Fig. 1.3: Fig. 1.3a shows the baseline V3, and Fig. 1.3b shows an image in flight, here without the KCU.

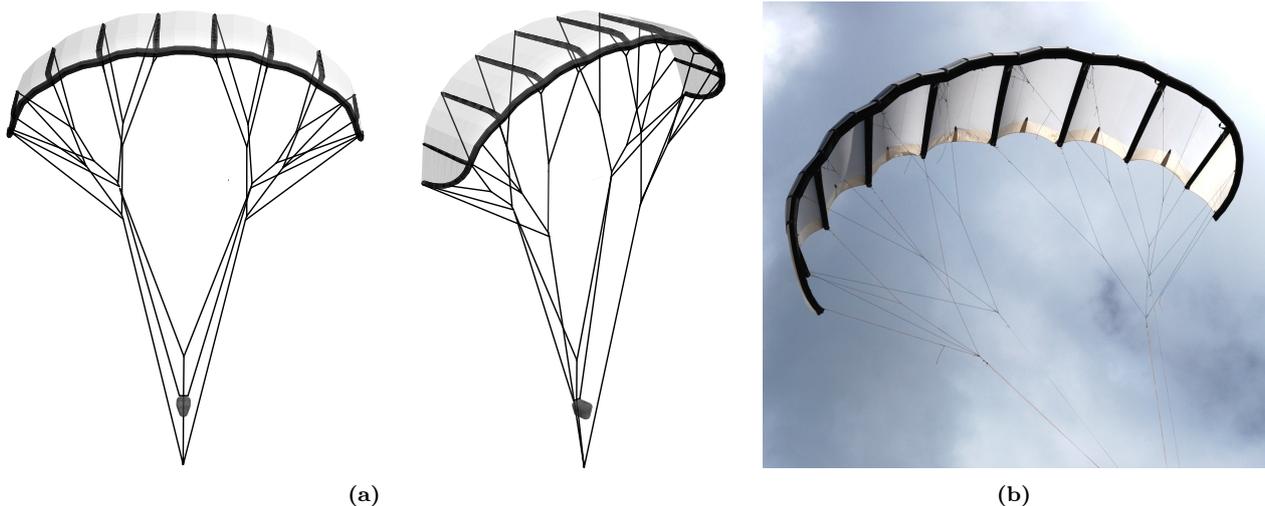


Figure 1.3: TU Delft V3 kite used in this study, with different bridle layout. (a) Two-view illustration of the baseline V3. (b) Experimental configuration without an airborne KCU [10].

Structurally, the V3 features an inflatable leading edge running spanwise and eight inflatable struts oriented chordwise. Internal pressure in this frame helps the wing retain its shape in flight. This inflatable structure makes operations and ground handling easier. A thin canopy, the primary aerodynamic structure, spans between the struts and the leading edge, while the trailing edge (TE) is not supported by an inflatable tube. Bridles are attached to both the leading edge (LE) and trailing edge; the LE bridle lines carry the majority of the aerodynamic loads, whereas the TE lines are primarily used for control, enabling depowering and steering inputs. The principal geometry parameters are summarised in Table 1.1. More information on the TU Delft V3 kite is available online¹.

Table 1.1: Design geometry of the V3 kite.

Property	Symbol	Value
Wing height	h	2.8 m
Projected wingspan (width)	w	8.3 m
Maximum chord	c_{\max}	2.7 m
Flattened surface area	S	25 m ²
Projected area	A	19.75 m ²
KCU mass	m_{KCU}	23.25 kg
Wing mass	m_{wing}	11.0 kg

1.3 Motivation

Airborne wind energy with soft LEI kites involves a tightly coupled aero-structural interaction: aerodynamic loads deform the wing, the deformed shape modifies the effective geometry, and this

¹https://github.com/awegroup/TUDELFT_V3_KITE

altered geometry feeds back into the aerodynamic loads over time. Modern models capture this coupling either by solving the fluid and structural sub-problems together (monolithic) or by iterating between them within each time step (partitioned), enforcing equilibrium of forces and kinematics on the interface until convergence [12]. For AWE design, reduced-order aero-structural models are attractive because they retain the essential two-way physics at practical runtimes. Typical pairings include a particle-system geometry for the inflated frame and bridles coupled to a fast aerodynamic solver - e.g. a 3D non-linear vortex-step method with viscous polars [13] — or simplified aerodynamic loading coupled to a particle system that reproduces symmetric depower and asymmetric steering deformations [10]. The virtual wind tunnel (VWT) framework solves the coupled aero-structural equilibrium of a soft-wing kite at prescribed operating conditions, and can be applied both to static configurations such as a kite held at zenith and to dynamic flight scenarios [14]. Building on this concept, a full flight is represented as a sequence of quasi-steady operating points, an approach that has recently been extended to steady circular cross-wind flight states.

Although modelling approaches for LEI kites are becoming increasingly sophisticated, reliable in-flight data to support them are still scarce, as no experimental method has been established. Only a few studies provide three-dimensional information on kite geometry under real pumping-cycle conditions, and most of these lack the accuracy or realism needed for direct comparison with models [15]. Fully coupled FSI remains particularly challenging for soft wings and, without validating measurements, offers limited practical benefit for aerodynamic characterisation [16]. In practice, design tools still face a trade-off between accuracy and computational speed, which means that developers continue to depend on intuition and iterative build-and-fly testing [13]. Validated aero-structural models could change this by enabling virtual prototyping and optimisation, cutting down on costly prototype loops and speeding up the design process [17, 13]. Addressing this gap requires in-flight measurements of the actual deformations experienced by a full-scale LEI kite, leading to the research objective of this thesis:

"To quantify the large-scale deformations of a leading-edge inflatable kite used in airborne wind energy through in-flight experimental measurements."

1.4 Report layout

A comprehensive review of existing knowledge is provided in Chapter 2. This chapter summarises the aerodynamics of LEI kites, typical deformation modes, existing modelling approaches, and previous experimental efforts. It further evaluates a wide range of potential measurement techniques for capturing large-scale shape changes. Building on these insights, the conceptual research design is laid out in Chapter 3. This chapter restates the research objective and formulates the research questions. Here, a trade-off of the measurement techniques is performed. The outcome of this analysis guides the selection of stereoscopic photogrammetry as the primary method, with ultra-wideband (UWB) ranging adopted as a complementary approach.

The stereoscopic photogrammetry setup is detailed in Chapter 4. This chapter covers the basic principles, hardware selection, camera calibration, marker design and placement, accuracy analysis and the full custom reconstruction algorithm. Chapter 5 introduces the UWB subsystem. It describes the underlying ranging principle, the Pozyx sensors used, their calibration, and an evaluation of achievable accuracy and communication range. Wind-speed and angular measurements are described in Chapter 6. This includes the Pitot-tube system, calibration of the apparent wind sensor, sensitivity analyses as well as angular vane design and their respective calibration. The full experimental configuration is presented in Chapter 7. This chapter describes how all subsystems—photogrammetry, UWB, wind-speed sensing, and KCU actuation—are integrated and synchronised. It further details the KCU inputs used for (de)powering and steering, the intrusiveness analysis, and the complete synchronisation chain from onboard sensors to UTC.

Chapter 8 outlines the post-processing workflow applied to all recorded signals. General filtering and correction steps are followed by photogrammetry-specific procedures, including correction for camera-rig deformation, transformation into a wing-fixed reference frame, outlier rejection, and wing-shape reconstruction. The experimental results are presented in Chapter 9. The reconstructed in-flight wing geometries are compared against the CAD reference, and deformation modes such as billowing, twist, depower-induced shape change, and steering-induced asymmetries are quantified.

Chapter 10 evaluates the consistency between measured deformations and predictions from a structural finite-element model (FEM). A static hanging-kite experiment is used as a validation case, and the resulting deformation fields are compared with FEM predictions using correlation metrics. Finally, Chapter 11 discusses conclusions of the research, while Chapter 12 provides recommendations for methodological improvements and directions for future work.

Literature review 2

Measuring the deformation of a fast, flexible membrane wing in real flight is intrinsically difficult. The kite operates at altitudes above 300 m and can see apparent wind speeds exceeding 35 ms^{-1} , so manoeuvres are rapid, load cycles are large, and the dynamic response ranges from slow global bending to higher-frequency local modes. The structure itself is an inflated fabric membrane with strongly coupled aero-structural behaviour: panels wrinkle, seams and bladders shift, and the overall shape adapts continuously to control inputs and turbulence. Instrumentation is constrained by mass, power, and robustness limits, while long standoff distances and safety restrictions limit direct access and make ground-to-air synchronisation and reliable telemetry non-trivial. Achieving high accuracy on meter-scale deformations under these conditions demands measurement approaches that are precise, time-aligned, and minimally intrusive.

This literature review establishes a focused basis for measuring these in-flight deformations of LEI kites. It concentrates on what will be used in this research rather than surveying techniques that proved impractical in the field. Specifically, the review (i) summarises the aerodynamics of LEI kites that drive and respond to structural deformation (Section 2.1); (ii) outlines current kite models with fluid–structure interaction (FSI) coupling (Section 2.2); (iii) describes typical, performance–relevant deformation modes (steering, depower, vibrational, Section 2.3); (iv) evaluates previous attempts at kite measurements and why accuracy has lagged (Section 2.4); and (v) examines practical techniques for this project, with emphasis on stereoscopic photogrammetry (global shape), UWB ranging (select purposes), and on–kite wind measurements (Section 2.5).

2.1 Aerodynamics of a LEI kite

LEI kites share the basic lift and drag mechanisms of lifting wings [18]. Therefore, faster flow along the suction side and slower flow along the pressure side is present [19, 20]. For LEI kites, the inflated leading edge (LE) acts as a bluff element that fundamentally shapes the near-field flow: computational fluid dynamics shows a sizeable pressure-side recirculation region just aft of the leading edge whose extent depends strongly on Reynolds number (Re) and angle of attack [21]. Beyond the dominant pressure-side bubble, two quasi-two-dimensional features are important. First, a laminar separation bubble forms on the suction side at moderate Re : the boundary layer lifts off shortly downstream of the LE, transitions within the shear layer, and reattaches further aft. In computations this appears as a short recirculating pocket whose length and location vary strongly with angle of attack and Re ; in the pressure distribution it produces a local plateau and a kink in the $C_p(x)$ curve prior to reattachment [21, 22]. Second, trailing edge (TE) separation can occur under adverse pressure gradients, especially in depowered states and when local membrane tension is low. Both effects interact with the overall deformation and are part of the flow topology in Fig. 2.1, where the separation and reattachment points along the chord are highlighted in the C_p – x distribution.

Three-dimensional effects are not a small correction but a defining feature of LEI aerodynamics. High-fidelity studies report spanwise transport that alternates between struts and reverses direction farther from the canopy, while chordwise vorticity grows towards the tips; introducing struts steepens spanwise gradients in velocity and vorticity [23, 24]. The span-wise flow around these struts generates more extreme gradients and absolute values in flow velocity and vorticity [25]. Increasing anhedral accelerates spanwise flow and modifies the wingtip vortex system, changing the induced-drag footprint in ways that couple directly to the kite’s deformation state.

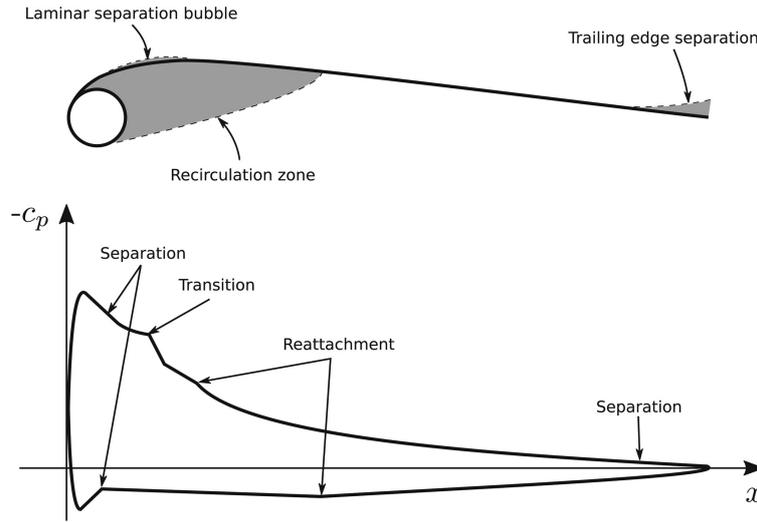
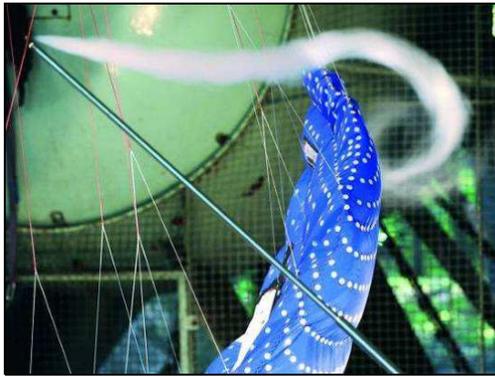


Figure 2.1: Flow topology around a LEI kite, showing pressure-side recirculation due to the inflated leading edge and a suction-side laminar separation bubble [21].

Membrane structures change both separation behaviour and spanwise load distribution compared with rigid airfoils. Permeable flexible wings have been investigated using particle image velocimetry (PIV), hot-wire anemometry (HWA), digital image correlation (DIC) and smoke visualisation, and reported three consistent trends: a suction-side separation bubble still forms and grows with freestream velocity and angle of attack; the recirculation region on the suction side is significantly smaller than for a comparable rigid wing because the membrane adapts its shape under load; and the tip-vortex system is strengthened, with an associated increase in lift due to the adaptive camber effect [26]. These findings highlight how deformation and aerodynamics are tightly coupled in fabric wings. On real kites the same features are observed: the left panel in Fig. 2.2a shows coherent tip vortices in smoke visualisation [27], while the right panel in Fig. 2.2b shows pressure-side separation on a LEI kite during manoeuvre, consistent with the pressure-side recirculation discussed above [17].



(a)



(b)

Figure 2.2: Visual evidence of three-dimensional flow structures on kites using smoke visualization techniques. (a) The tips show coherent tip vortices on a RAM-air kite [27]. (b) Pressure-side separation is visible on a LEI kite [17].

Across typical operating conditions the relevant Re span orders of magnitude. For a pumping cycle, Re in the order of 10^5 are representative for the retraction phase of AWE systems, whilst Re in the order from 10^6 to 10^8 are more representative of the traction phase [21]. This range implies that both the suction-side laminar separation bubble and the pressure-side recirculation region are sensitive to operating conditions and deformation, and vice versa. In some cases, it might therefore be beneficial to use trip strips, transitioning the flow [28]. It follows that interpreting performance reliably requires wing shape and inflow to be measured together whenever possible.

The aerodynamic features described above are tightly coupled to the kite’s structural response, and the deformations in turn reshape the flow. In summary the inflated membrane adapts its local camber and curvature under load, which modifies boundary-layer behaviour and pressure gradients compared to a rigid wing: the suction-side bubble can shorten as the membrane adapts, while the pressure-side recirculation behind the leading edge grows or shrinks with the instantaneous shape. Three-dimensional transport driven by anhedral and the presence of struts redistributes momentum spanwise and alters chordwise vorticity towards the tips, so that even modest geometric changes can shift separation and reattachment locations.

2.2 Modelling a kite

A kite is different from a free wing: the tether introduces a large external force and a geometric constraint that changes both the equilibrium and the dynamics. In the simplest static case (i.e. parking the kite in Fig. 2.3), the aerodynamic forces L and D , the weight mg , and the tether force F_T must balance at a given elevation β for a specified wind speed \mathbf{v}_w [29]. This immediately motivates modelling choices that differ from conventional aircraft wings due to the presence of a tether.

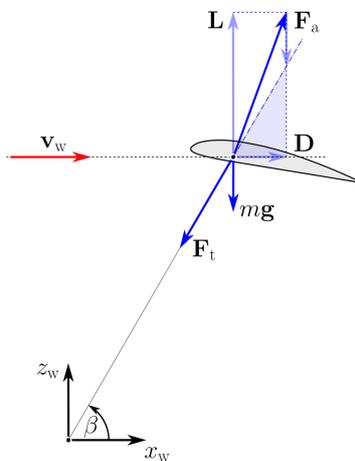


Figure 2.3: Force equilibrium in static flight [29]. Here F_T is the tether force, mg the weight of the kite, L and D the lift and drag, \mathbf{v}_w the wind speed vector and β the kite elevation angle.

A convenient description of the flight path of a tethered kite is through spherical position coordinates (β, ϕ, r) together with a local tether frame (x_t, y_t, z_t) aligned with the kite heading and the tangent plane of the motion sphere, and a heading angle ψ that relates the kite’s longitudinal axis to the meridians (Fig. 2.4) [16].

With the force balance and kinematics fixed, the modelling task is to predict how the deformable structure and the aerodynamics co-evolve for given inputs (steering, depower) and flight states. Below this logic is followed: first define model fidelity, then specify the structural discretisation, then detail the aerodynamic coupling on the deformed geometry, and finally state the measurable outputs used for validation.

Levels of model fidelity

A widely used low-order structural representation is a particle system whose nodes are connected by linear or nonlinear spring–damper elements [30]. Aerodynamics can then be modelled by a vortex step method (VSM) defined on a segmented lifting surface that follows the instantaneous structural shape. The resulting coupled model advances the structure under aerodynamic and inertial loads while updating the aerodynamic solution on the deformed geometry at each time step. This particle system model (PSM) and VSM combination has been applied to soft-wing kites to study actuation-induced and aeroelastic deformations in a quasi-steady “VWT” setting [10, 14]. In

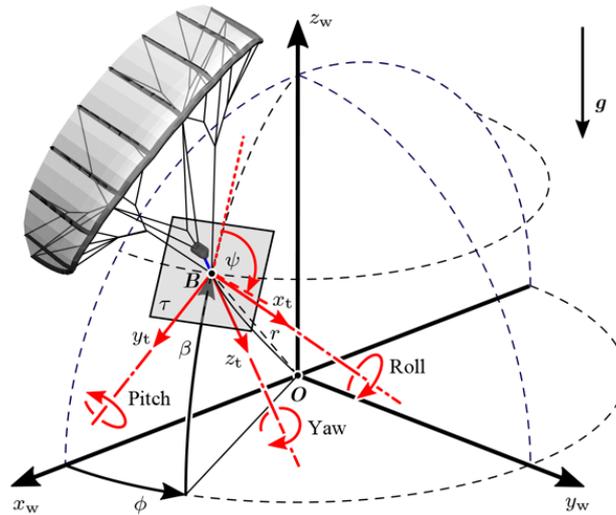


Figure 2.4: Coordinate systems for a tethered kite: ground frame, spherical position (β, ϕ, r) , local tether frame (x_t, y_t, z_t) , and heading ψ [16].

that framework, actuation (steering and depower) alters geometric constraints at bridle attachment nodes, and the membrane shape adjusts until the iterative process converges. At the highest fidelity, CFD–FEM resolves viscous, separated flow and large membrane deformations directly, increasing accuracy at far higher costs than the PSM–VSM surrogate.

Structural model

In the PSM, node locations are chosen to coincide with physically meaningful points: bridle–to–canopy attachments at the leading and trailing edges of each strut and bridle junctions. Springs and dampers represent in–plane stiffness and energy dissipation; gravity and inertial terms are included on each mass. Aerodynamic loads are applied at, or distributed between, these nodes according to the VSM pressure field (Fig. 2.5) [10]. This discretisation fixes what an experiment should observe for validation: three–dimensional time series of the nodal positions, from which spanwise deflection, local twist and key baselines (tip–tip, strut–strut) can be computed. Performing measurements at these nodes minimises interpolation error when comparing model and data.

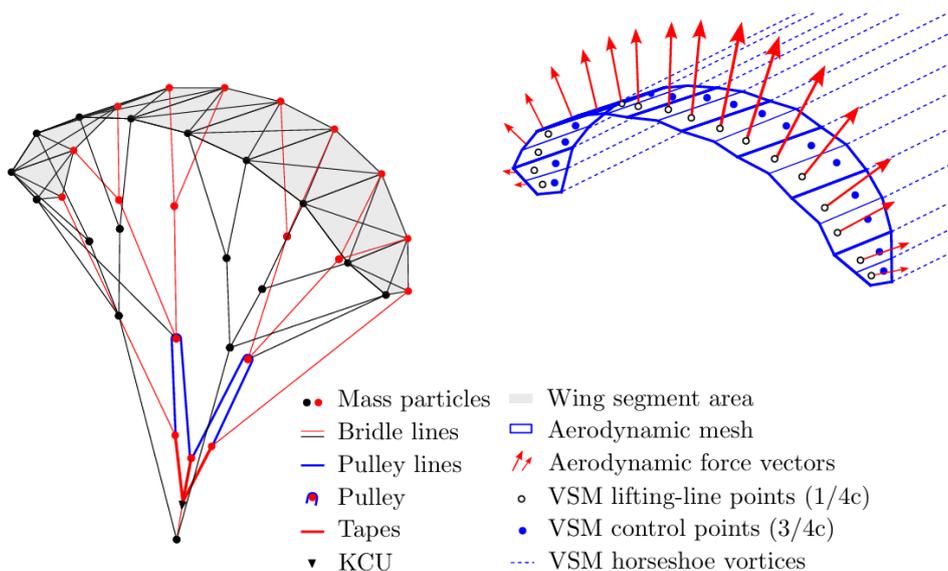


Figure 2.5: Coupled aero–structural model of the TU Delft V3 kite [10]. Left: distributed mass–spring–damper network for the structure, with nodes placed at bridle attachment and structural junctions. Right: vortex step method on a segmented lifting surface conforming to the deformed shape.

Aerodynamic model

The VSM treats the deformed kite as a set of bound and shed vortices on a faceted surface. Loads depend on the local incidence and on the instantaneous geometry provided by the structure. In the quasi-steady formulation, each input or flight state is treated as an equilibrium in a VWT, which is appropriate when the dominant deformation content is at low frequency. Reported deformation spectra for LEI kites indicate that most global deformations occur below roughly 2 Hz, supporting quasi-static assumptions for many manoeuvres [31, 14]. When rapid transients are of interest, the PSM can still be run in the time domain with structural damping and inertia, or augment the aerodynamics with unsteady corrections if required, while remaining far cheaper than full CFD-FEM.

A simple baseline for intuition is a flat-plate equilibrium model in which lift, drag, mass and tether force act about a bridle point; equilibrium requires vanishing moments and a consistent relation between heading, elevation and lift-to-drag ratio [32]. Such models are useful for guidance and checks but are not sufficient for deforming membranes and are therefore complemented by the coupled aero-structural approach above.

Kinematics, inputs and outputs

Expressing the coupled aero-structural solution in the tether frame of Fig. 2.4 enables direct comparison with measured trajectories and control inputs defined on the bridle system [16]. Steering and depower commands are imposed as changes in line lengths or attachment geometry, which act as inputs to the coupled aero-structural solver. The solver computes the resulting nodal positions, velocities, and internal loads of the kite and bridle network. These outputs are post-processed to obtain aerodynamically relevant deformation metrics, which form the basis for comparison with experimental data and define the targets for model validation.

Scope and limitations

The PSM-VSM approach reproduces the large scale deformations needed for performance analysis while keeping runtimes practical [10, 14]. Its accuracy depends on appropriate placement of nodes, realistic stiffness and damping parameters, and consistent aerodynamic calibration on the deformed geometry. Assumptions of quasi-steady flight and idealised manoeuvre paths ease interpretation but introduce bias when strong unsteadiness or asymmetry is present; these cases require time-resolved coupling.

Taken together, these elements define a clear contrast between simulation and experiment: the model expects node-level shape and load or apparent wind speed information in the tether frame, and the measurement strategy should be designed to supply it at the required spatial and temporal resolution.

2.3 Typical deformations

Kites used for AWE are designed to deform. That flexibility enables steering and depowering, but it also couples the flow and the structure so that small geometric changes can have large aerodynamic effects. Steering is achieved by applying asymmetric tension to the bridle system, effectively shortening lines on one side while lengthening them on the other, which introduces asymmetric twist and warp and generates a yawing moment [33]. For power-cycling, it is beneficial to transition between powered and depowered states. Depowering reduces the angle of attack to lower aerodynamic forces during retraction, while powering restores lift and traction for the reel-out phase.

Kites must remain lightweight to enable operation in low wind speeds, which is achieved by minimising material usage and material selection [34]. As a consequence, structural stiffness is limited and aerodynamic loads are redistributed through a flexible load-bearing system consisting of the membrane, the inflatable structure, and the bridle network. The approach improves robustness

during development testing, where hard landings are common, because compliant structures distribute loads and mitigate failure [35]. The pressure in the canopy and the tensile bridle system provide the effective stiffness that sets the deformation field under load [34].

2.3.1 Depowering deformation

Depowering reduces the angle of attack and redistributes tension in the canopy and bridles. In practice the wing adopts a more curved planform with increased anhedral, and the projected area seen by the inflow decreases; both effects contribute to reduced aerodynamic forces during retraction [8, 34]. The characteristic geometry change between powered and depowered is shown in Fig. 2.6. For the flow, a lower angle of attack weakens suction, and changes the pressure-side recirculation region behind the inflated leading edge; the balance depends on the operating Re and the instantaneous shape, but the global result is a marked reduction in tether force during depower [8]. Though these shape changes have been qualitatively analysed, quantitative analysis is lacking.

2.3.2 Steering-induced deformations

Turning deformations are dominated by asymmetric twist and warp of the canopy. When the steering line set at one tip is shortened, the local angle of attack and the effective area on that side increase, while the opposite side reduces; the resulting lift and drag imbalance produces a yaw and roll moment, resulting in a turn [17]. The qualitative shape change is illustrated in Fig. 2.7. The turn rate depends on the steering command, apparent wind and kite-specific properties, and is commonly described by a turn-rate law that captures these dependencies [36, 37]. Aerodynamically, the pulled side develops stronger bound circulation and a more energetic tip vortex, while the slack side weakens, which modifies induced drag and lateral force during the manoeuvre [17, 36].

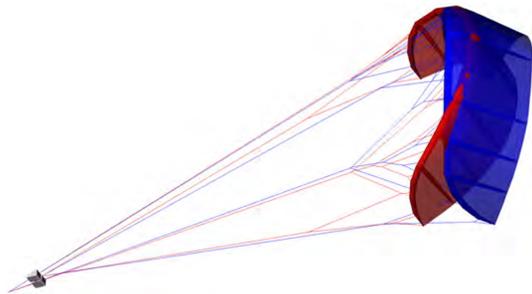


Figure 2.6: Depowering: shape of powered (red) and depowered (blue) kite [8].

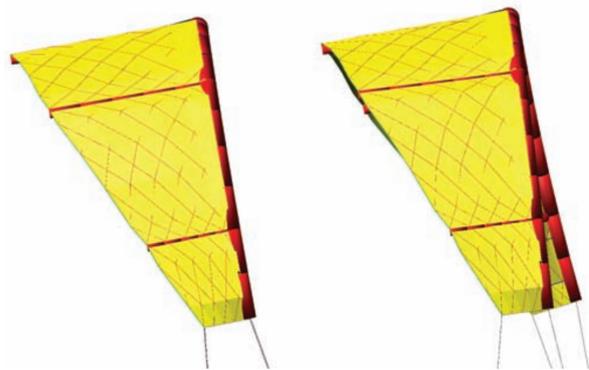


Figure 2.7: Steering: shape before (left) and after (right) a steering input [17].

Control inputs and their definition

Definitions in the literature sometimes use a bar-type control system with a depower trim and differential tip lines, as in [38]. Alternatively, in AWE a KCU is often used. The same inputs can be expressed in terms of line-length changes implemented by the KCU: a scalar power or depower setting u_p or u_{dp} shortens or lengthens the front versus back line group to change the angle of attack, while a steering input u_s creates a differential between left and right steering line groups to introduce asymmetric twist. This description is equivalent in purpose to the bar definitions once mapped to line-length differentials, and it is consistent with AWE implementations where the KCU realises these geometry changes [8, 38]. For reference, Fig. 2.8 shows the bar-based definition used in towing tests by Elfert et al. [38].

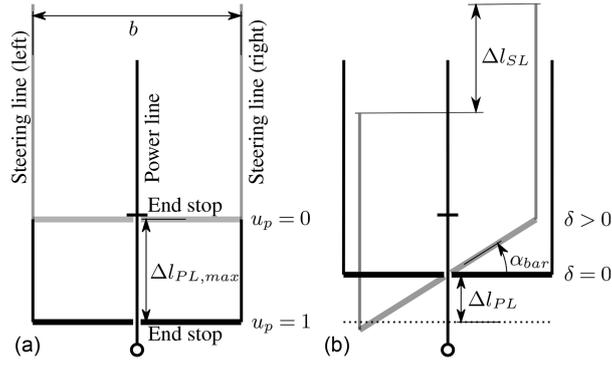
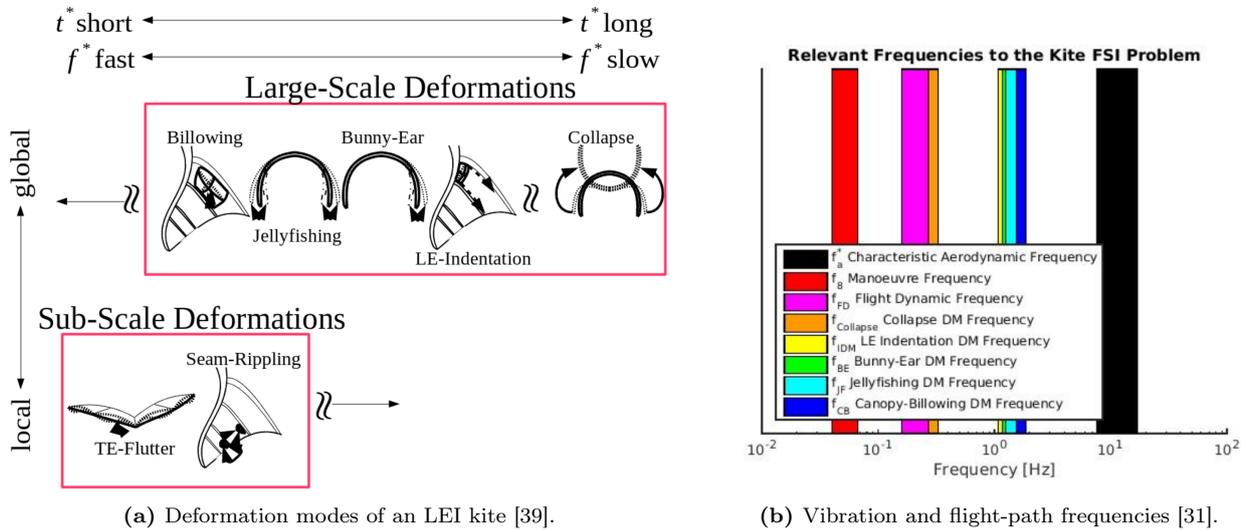


Figure 2.8: Common bar-based definitions of (a) depower and (b) steering input [38]. In a KCU, the same inputs are implemented as front–back and left–right line-length changes. Note that steering input in the image is defined as δ , whilst u_s is used throughout this report. Additionally, the power setting u_p is different from depower setting u_{dp} used in this research.

2.3.3 Vibrational deformations

Vibrational deformation modes in LEI kites can be grouped into global and local classes (see Fig. 2.9a)[39, 31]. For aeroelastic model validation the global modes are most relevant because these are simulated. Sub-scale deformations occur at part of the structure, and don’t provide an overview of the shape of the entire kite. Billowing (adaptive camber) is the baseline response under load, where the canopy inflates and redistributes camber. Jellyfishing denotes an oscillatory motion in which the canopy alternately flattens and tightens, often during cornering manoeuvres [36]. Bunny-ear flapping is a related asymmetric case where one tip oscillates strongly, typically in tight turns. Leading-edge indentation occurs at low angle of attack when the nose locally buckles inward and is a precursor to partial or full collapse.



(a) Deformation modes of an LEI kite [39].

(b) Vibration and flight-path frequencies [31].

Figure 2.9: Global and local vibration modes of LEI kites and their characteristic time scales. (a) Mode taxonomy highlighting billowing, jellyfishing, bunny-ear flapping and leading-edge indentation [39]. (b) Comparison of deformation and flight-path frequencies indicating that most global deformation energy lies below about 2 Hz [31].

To place these modes on a time scale, [31] compared deformation frequencies with flight-path dynamics (Fig. 2.9b). Most global deformation energy occurs below about 2 Hz, which supports quasi-static structural assumptions for many manoeuvres. By the Nyquist–Shannon theorem [40, 41] this implies a minimum sampling rate of 4 Hz, though higher rates are advisable to capture transients and maintain synchronisation with other signals.

Overall, typical deformations can be summarised by a small set of observables that link directly to aerodynamics: anhedral, projected area and span between powered and depowered states; spanwise deflection and local twist during steering; and the presence or onset of adverse modes such as leading-edge indentation. These are the quantities targeted by the measurement and modelling workflow in the remainder of this thesis.

2.4 Previous experimental kite studies

A range of studies has addressed the problem of reconstructing soft-wing kite shape, from controlled facilities to in-flight measurement on operational systems. Below, the principal contributions are grouped by setting and sensing approach, with emphasis on what was measured, what accuracy was achieved, and which constraints proved limiting for outdoor flight.

2.4.1 Wind-tunnel campaigns and controlled experiments

de Wachter [27]: Executed comprehensive wind-tunnel campaigns on a ram-air wing, capturing its 3D shape using multi-camera photogrammetry and a phase-difference laser scanner, measuring loads with calibrated load cells, and visualising boundary-layer behaviour via smoke and infrared thermography. The thesis contrasts the two shape-capture methods (photogrammetry proving superior for surface fidelity) and couples measured geometries to CFD to interpret aerodynamic impacts. These are high-quality, well-controlled experiments that quantify ballooning, leading-/trailing-edge deformations, and spanwise camber under set conditions. What remains missing is continuous, outdoor in-flight shape reconstruction at operational scales: the study is not representative for a kite in flight, manoeuvring through a pumping cycle.

Breukels [17]: Developed an engineering methodology for LEI kite design that combined in-flight documentation of steering-induced deformation with controlled wind-tunnel investigations, as well as simulations. The thesis reports qualitative and quantitative observations of asymmetric inputs producing twist/warp and associated shape changes, and includes wind-tunnel flow visualisation (smoke) to study separation behaviour on the pressure side of the inflated leading edge. However, the work does not present a full-field, multi-camera photogrammetric reconstruction of the canopy in flight. As a result, spatial coverage in outdoor experiments was sparse (few tracked features, monocular foreshortening), and the wind-tunnel validations operate at lower Re , short range, and without full tether dynamics or atmospheric variability, limiting direct transfer to operational AWE conditions.

Perin et al. [42]: Conducted a controlled wind-tunnel campaign on a scaled ram-air kite to validate a coupled FSI model. Experimentally, they measured lift, drag and line loads over different wind speeds and mapped the inflated canopy using stereo DIC on the lower surface (speckled fabric), achieving sub-millimetre geometric resolution. They also used qualitative flow visualization and compared measured shapes with LS-DYNA simulations. Despite the rich dataset, the method is intrinsically tunnel-bound: DIC requires extensive surface preparation, stable multi-view geometry and short stand-off distances; occlusions from lines/folds produced holes in the reconstruction; only the underside was captured, and time-resolved full-field outdoor tracking was not attempted. The work is therefore excellent for model calibration in controlled conditions, but it does not demonstrate centimetre-accurate, in-flight shape reconstruction.

Elfert et al. [38]: A tow-test procedure for an LEI kite was executed with a kite-mounted measurement suite: an onboard sensor system recorded position and orientation, and a low-cost multi-hole probe measured the relative flow-velocity vector (apparent wind). From these time-synchronised signals the steering behaviour was identified (including steering gain and dead-time). The campaign cleanly isolates input-response relations under repeatable conditions and is valuable for calibration of low-order turn-rate models. However, no in-flight shape (3D geometry) was reconstructed, and the

setup does not address the long stand-off distances or atmospheric variability of full-scale outdoor AWES.

Schwoll [43]: Presented a case study in which a V2 Mutiny LEI kite (including the full bridle) was modelled in MADYMO using a uniform-pressure FEM. To isolate the structural response from uncertain aerodynamics, the kite was hung upside-down indoors and subjected to three static load cases: gravity only, a concentrated load at the leading-edge centre, and a falling-mass impact. The bridle geometry and inflation pressure were included in the model, and qualitative deformation comparisons were made between simulation and photographs of the suspended kite under each load. This approach cleanly removes wind/tether effects and captures the bridle’s influence on stiffness, but it does not deliver a time-resolved, full-field 3D shape reconstruction; comparisons remain qualitative and limited to static cases, with the model noted to be stiffer than the real kite.

2.4.2 In-flight visual tracking

Hanke & Schenk [44]: Demonstrated single-camera, platform-mounted photogrammetry on a paraglider in flight by tracking marked canopy points and reconstructing 3D geometry from multi-image views, their setup depicted in Fig. 2.10. The study shows that monocular 3D is feasible when calibration is stable and markers stay visible, but frequent motion blur, occlusions, and attitude changes interrupt tracks; absolute accuracy is sensitive to calibration drift and out-of-plane motion accuracy is particularly low. The setup does not represent AWE length scales or ranges and does not deliver continuous, full-field canopy shape over aggressive manoeuvres, leaving gaps for outdoor, centimetre-level deformation reconstruction. Their setup also required a pilot to be flying with the glider, which would be incredibly dangerous if not impossible for an AWES kite.

Behrel [45]: Using synchronised onboard video and kinematic logs (IMU/GNSS), Behrel contrasted canopy shape in quasi-straight flight (rotation rate $\approx 0.1 \text{ rads}^{-1}$) versus a high-rate turn ($\approx 3.9 \text{ rads}^{-1}$), and illustrated the structural deformations through these turns. The study shows increased anhedral/tip deflection and overall billow during the turn relative to straight flight—i.e., stronger steering correlates with larger structural deformation. The observation is qualitative: a single-camera view and navigation data were used, without multi-camera photogrammetry, so no centimetre-accurate 3D shape reconstruction was provided.

Poland & Schmehl [10, 34]: In-flight video from a camera on the KCU was used for geometry by extracting stills at extreme states (powered reel-out, depowered reel-in) and correcting fisheye distortion; because the camera-kite distance varied with KCU swing and kite pitch, absolute scale was avoided and measurements were expressed as percentage changes then back-scaled with the design geometry. Strut lengths were assumed constant to define local metre-per-pixel ratios, enabling

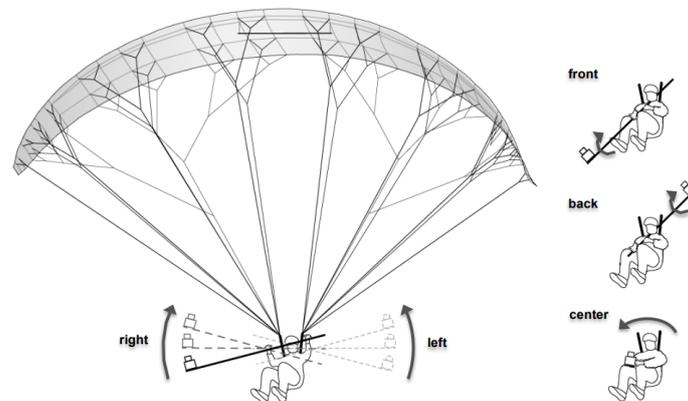


Figure 2.10: Single-camera photogrammetry setup used by Hanke & Schenk in flight. The setup is one of the few visual tracking setups used whilst flying.

estimates of span change and TE strut tip–tip “billowing.” Turning footage was deemed too distorted for quantitative use. Using this workflow, the powered–depowered comparison showed a decrease in projected width ($\sigma \sim 3\%$) and an increase in TE strut tip–tip distance indicative of billowing. Methodological limitations include unknown in-flight extrinsics, reliance on stills rather than time-series, the constant-strut assumption, and lack of reference data for absolute scale. Therefore, the results are strongest as relative changes rather than absolute shape, and still lack accuracy.

Leuthold [39, 31]: Used flight video of LEI kites to identify vibrational deformation modes and estimate their characteristic frequencies by frame counting (order-of-magnitude timing from footage). This yields a qualitative taxonomy and time scales for global modes but no multi-camera 3D shape, strain fields, or accurate geometry, meaning quantitative deformation data are largely absent.

Infanzon [46]: Investigated reconstruction of a surf–kite’s flying shape using consumer photogrammetry. The setup placed two (ultimately three) GoPro Hero3 Black cameras on a rigid ground rig, synchronised via the GoPro Wi-Fi remote, and processed stills in Agisoft PhotoScan Professional; frame matching was checked post-hoc with a visible stopwatch to ensure epoch consistency. Camera intrinsics were calibrated with a chequerboard; field-of-view and expected precision were analysed, with measured horizontal and pixel size. Early boat-tow trials produced unusable footage; the successful campaign fixed the kite to a ground anchor and captured short, synchronised bursts, from which discrete 3D shapes and simple trajectory parameters were recovered. However, the method depends on a stationary, long-baseline ground rig, where conditions are much more controlled than in-flight. In flight a setup would be more complex, as intrusiveness plays a role, and the length from ground-based cameras to the kite are limited.

Costa [47]: Performed tow-tests with an LEI kite using a stereoscopic, two-camera rig mounted on a beam on a towing vehicle and short lines to keep the canopy within the calibrated field of view. The kite was covered with randomly patterned targets and processed with 3D DIC to obtain time-resolved displacements; ground sensors (tether load, wind, angles) were logged on the same PC to align timestamps, and rigid-body motion was removed by registration before analysing deformations versus load. This produced spanwise deformation trends and qualitative coupling to aerodynamic loading. The research does not measure a kite representable of an AWES, and the stereoscopic rig is positioned on the truck, where there are less restrictions on mass and size.

2.4.3 Distributed kinematic/structural sensing on the wing

van der Vlugt et al. [8]: Instrumented a pumping-kite system with an Xsens MTi-G IMU on the centre strut, GNSS on the kite/KCU/ground station, a tether load cell and a bridle-mounted Pitot for inflow; data were fused (EKF) to estimate the kite’s kinematics and turning behaviour in field tests. The campaign quantified system-level dynamics but did not reconstruct the canopy shape; GNSS accuracy and dropouts limited fine-scale geometry and local deformation inference.

Jonard [15]: Deployed a line of lightweight MEMS IMUs along the leading edge to estimate in-flight shape, and used ground-based two-camera photogrammetry for calibration/validation during controlled manoeuvres. IMUs provided continuous orientation histories, but absolute heading and drift were sensitive to mounting/alignment and magnetic disturbances. The low amount of sensors required a large portions of the wing to be reconstructed by an estimated wireframe, as depicted in Fig. 2.11. The stereo setup suffered from marker visibility and synchronisation limits. Due to the large distance to the kite - together with insignificant calibration procedures - the accuracy of the stereoscopic setup was extremely low. As a result, full-span, centimetre-level 3D reconstruction was intermittent outdoors.

van Reijen [48]: Attempted in-flight sensing on a ground-fixed LEI kite using IMUs at the leading-edge centre and both tips to extend shape measurement to large kites and turning manoeuvres.

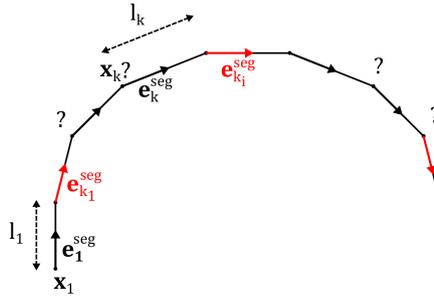


Figure 2.11: Wireframe model as described by Jonard [15]

Measurements were corrupted by magnetic interference from an onboard battery system, yielding unreliable headings and degraded attitude; consequently, no robust 3D shape reconstruction was obtained.

Schelbergen & Schmehl [33]: Flight tests on the V3 LEI kite instrumented the wing/KCU with two Pixhawk units and a bridle-mounted multihole Pitot with AoA vane to log apparent wind, attitude and kinematics; the dataset was used to analyse turning and KCU swinging dynamics. The study delivers high-quality in-flight kinematic/airdata under operational loads but does not reconstruct the canopy’s 3D shape; absolute accuracy remains tied to onboard calibration (no sideslip sensor) and moving-platform extrinsics.

2.4.4 Wind-speed measurements

Oehler [49]: Designed, calibrated and flight-tested an airborne air-data boom mounted in the bridles above the KCU (see Fig. 2.12). The setup combined a Pitot-static tube and two vane encoders to measure V_a , α and β_s in situ, with wind-tunnel calibration and error analysis for pressure/temperature offsets and yaw sensitivity. Using synchronised logs (GPS/IMU/KCU), the work derived the resultant aerodynamic coefficient c_R and showed that it depends strongly on wing loading and, to a lesser extent, power setting; the kite’s angle of attack (AoA) varied only within a narrow band, with high-frequency α oscillations correlating with reel-speed fluctuations. A geometric depower model linked tape length to depower angle. Limitations: the method does not reconstruct canopy geometry and assumes near-free-stream conditions several metres below the wing; nevertheless, it establishes a practical framework for pairing air-data with kinematics for model validation.

Oehler & Schmehl [16]: Extended the approach to full pumping-cycle characterization using a rigid frame on the power lines carrying a self-aligning Pitot and two vanes, synchronised with KCU/IMU/GNSS. The study reports substantial variations of aerodynamic coefficients with trim (power setting), steering and flight direction, while angle of attack remains in a narrow range; representing coefficients as functions of angle of attack alone is inadequate, highlighting aeroelastic deformation effects. During constant-force crosswind manoeuvres, angle of attack is inversely related to V_a . The method provides operational-scale reference data for L/D and coefficient trends; like the 2017 setup [49] it does not measure 3D shape, so it complements (rather than replaces) other measurements.

Borobia-Moreno et al. [50]: Employed an onboard multi-hole air-data system (rigid boom on the kite) to measure true airspeed (TAS), AoA and sideslip (AoS) directly; the calibrated sensor provided accuracies of about $\pm 0.5 \text{ ms}^{-1}$ for TAS and $\pm 1^\circ$ for AoA/AoS. Air-data were synchronised with IMU/GNSS, magnetometer, tether load cells and a ground wind station, then fused in a continuous-discrete EKF (estimation-before-modelling) to improve state estimation and identify aerodynamic coefficients. LEI tests were constrained by short tethers and low winds; AoA frequently exceeded the instrument range ($> 50^\circ$, post-stall), limiting usable records, whereas a rigid delta

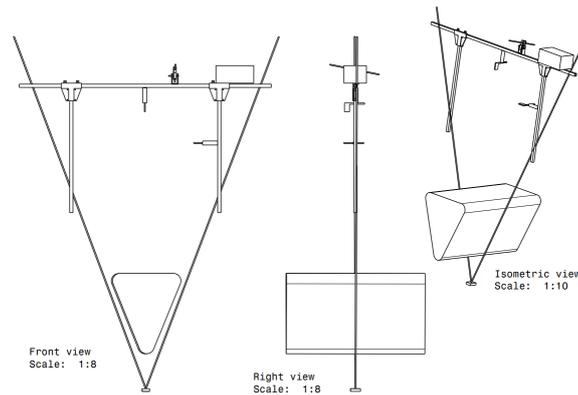


Figure 2.12: Bar construction designed by Oehler [49] to carry the Pitot tube and other measurement units, and perform wind speed measurements on the flying kite.

kite yielded repeatable figure-eight data. The setup markedly improves apparent-wind estimation for aerodynamic identification but does not reconstruct wing geometry, and residual uncertainty in position and angular rates remains due to system limitations.

Takeaways

Across the record, accurate kite-shape measurements split into two camps. In controlled facilities (wind tunnels, towing rigs) full-field geometry has been recovered with multi-camera photogrammetry or DIC and linked to loads, but these data do not represent outdoor AWES conditions with long tethers, high apparent wind and aggressive manoeuvres. In flight, visual approaches (monocular KCU video, ground-based stereo, single-camera paraglider studies) have yielded qualitative or intermittent reconstructions and relative shape changes, yet absolute, centimetre-level accuracy degrades with range, occlusion and changing viewpoints, and continuous full-span coverage through turns is not maintained. Distributed kinematic sensing (IMU/GNSS/Pitot on the kite/KCU) has delivered robust airdata, kinematics and steering laws, but not the 3D canopy geometry; IMU arrays on the leading edge help with global modes (bending/twist) but suffer from heading drift, alignment sensitivity and sparse-to-shape ambiguity. Consequently, accurate full kite 3D shape reconstruction in-flight on a kite used in AWES remains challenging. Kite shape reconstruction has only been done in controlled conditions or without sufficient accuracy.

2.5 Measurement techniques investigated for kite deformation measurements

This section will go into the different measurements researched for kite deformation studies. Though more techniques were researched thoroughly, the focus is on photogrammetry and UWB, as these are the techniques used in this study. This is substantiated in Chapter 3.

2.5.1 Photogrammetry

Photogrammetry is a visual measurement technique that reconstructs 3D geometry from 2D images. In stereoscopic setups, triangulation is performed by projecting rays from synchronised cameras through image points, with the intersection yielding the 3D position of markers [51]. Compared to alternatives such as DIC [42], which requires dense speckle patterns, photogrammetry could rely discrete high-contrast markers, making it more suitable for large, flexible structures such as kites.

Marker design and placement are critical for accuracy. It has been reported that although inertial sensors provided smoother signals, photogrammetry was essential for correcting drift [15]. However, the same study showed that small markers became unrecognisable during turns, highlighting the need for large, high-contrast designs [15]. This was a result of the camera being on the ground,

with the setup described in Fig. 2.13. Wide chequerboard crosses with a well-defined centre for outdoor visibility are generally recommended. In wind tunnel experiments, the potential of large multi-camera systems was demonstrated by reconstructing the full 3D geometry of a RAM-air kite with 14 cameras and 2000 markers [27]. While highly accurate, such setups are impractical outdoors due to cost, complexity, or camera to object distance.

In-flight applications have been attempted at smaller scales. Accurate 3D tracking has been shown to be possible with a single calibrated camera on a telescopic rig, provided careful calibration and marker distribution are used [44]. Tip deformations during turning manoeuvres have been analysed using photogrammetry with fixed reference points, showing that out-of-plane motion can be resolved if reference geometry is well defined [45]. More recently, single-camera photogrammetry from the KCU has been applied to compare powered and depowered kite states, though accuracy was limited by the unknown camera orientation [10].

Beyond these controlled or single-camera cases, there are a few stereoscopic field attempts under less controlled conditions. One study used a car-towed stereoscopic rig with a fixed baseline and an image-correlation workflow to capture canopy deformations during straight-line runs; while occlusions and short stand-off distances limited full-field coverage, the approach demonstrated feasible outdoor shape tracking and linked deformation to load changes [47].

Another study mounted multiple action cameras on a stationary ground rig in local winds, synchronised them wirelessly at millisecond-level timing, and reconstructed the 3D flying shape via close-range photogrammetry; practical accuracy estimates indicated in-plane precision on the order of centimetres and depth precision on the order of several centimetres at typical baselines/line lengths, improving with shorter lines or larger baselines [46].

Accuracy strongly depends on the setup. Mean errors of 1.32 mm have been reported in close-range single-camera tests [52], whereas stereoscopic systems have demonstrated sub-millimetre accuracy in laboratory conditions [53]. For field experiments, the main limitation is the reduced number of viewpoints, motion blur, and variable lighting conditions. Re-identification methods for objects leaving the field of view have been proposed [54], while background subtraction challenges in dynamic outdoor environments (e.g., moving clouds) have been reviewed [55]. Despite discussed limitations, stereoscopic photogrammetry remains among the most precise and practical techniques for measuring kite deformations.

Modern camera technology further enhances feasibility. High-resolution GoPro systems allow 5K video at 30 Hz, with higher frame rates at lower resolutions. These lightweight, rugged devices can be mounted directly on the kite or support structure, and have been shown to withstand crashes

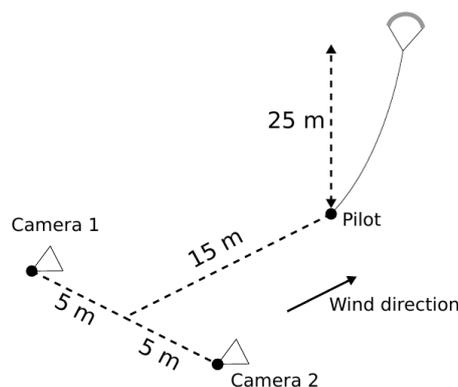


Figure 2.13: Most recent research into kite deformations using stereoscopic photogrammetry: the experimental ground stereoscopic photogrammetry setup used in [15].

and high accelerations without failure. With appropriate marker design and calibration, stereoscopic photogrammetry can achieve centimetre- to millimetre-level accuracy at kite stand-off distances, making it a suitable optical method for in-flight kite deformation measurements.

2.5.2 Ultra-wideband modules (UWB)

Ultra-wideband (UWB) is a short-range wireless communication and ranging technology that transmits extremely short radiofrequency pulses over a wide frequency spectrum. Distance is derived from the time of flight of these pulses, offering higher resolution than acoustic or radar-based methods due to the speed of light and large bandwidth [56, 57]. In practice, UWB can achieve direct distance measurements between nodes with errors of only a few centimetres.

Reported accuracy varies with system configuration. Under optimal indoor conditions, UWB modules achieve millimetre-level precision [58], but in simpler single-anchor ranging setups, errors of 2–10 cm are more typical. Grosswindhager et al. [59] developed the SALMA single-anchor localisation system, showing mean errors of 8 cm, with 99% of errors below 30 cm. Accuracy was found to improve outdoors, where multipath interference is less severe. Polonelli et al. [60] further demonstrated that orientation of the sensors strongly influences precision, reporting errors below 5 cm at short ranges but increasing rapidly with distance. Line-of-sight conditions are essential, as non-line-of-sight propagation significantly degrades performance.

The temporal resolution of UWB is a major advantage. With pulse durations as short as 0.16 ns at 500 MHz bandwidth, ranging can be performed at high sampling rates (for more recent sensors), enabling dynamic tracking of structural motion [57]. Filtering techniques, such as Kalman filters, have been shown to improve distance estimates in noisy conditions [59, 61], although achieving millimetre-level accuracy remains unrealistic without multi-anchor setups. Liu and Bao [61] highlight that accuracy can be improved at the cost of significant computational and memory resources, which may be impractical for real-time outdoor tests.

For kite deformation measurement, the main challenge lies in scale. Multi-anchor localisation is not feasible at flight altitudes of several hundred metres or at fixed positions mounted on the flying kite, restricting applications to single-anchor systems. While this limits spatial resolution, UWB could still provide centimetre-level accuracy for measuring distances between nodes, such as kite tips. Errors due to receiver delay can be quantified and corrected by calibration against known distances, while averaging consecutive samples or applying filters can further suppress noise. Multipath reflections remain a concern, but these can be minimised with careful placement of modules. Despite limitations, UWB offers an attractive trade-off between accuracy, cost, and robustness. The modules are compact, lightweight, and inexpensive. With proper calibration and filtering, single-anchor UWB ranging can reliably provide centimetre accuracy in dynamic conditions, making it a promising complementary technique to optical methods for capturing global deformation modes.

2.5.3 Others

This subsection briefly summarises other measurement techniques that were surveyed but not adopted in the present work. They are included for completeness and to record key takeaways for the later trade-off; detailed descriptions, datasets and methodological notes are compiled in a separate, extended literature survey. At times a single statement is supported by multiple citations: this is deliberate, to include all literature without extensive unnecessary explanations of techniques not implemented in the current study. Therefore, closely related findings from several sources are merged, with the references retained to substantiate each claim.

Apart from the techniques stated below, a reconstruction through load cells in the bridles was considered. However, this can infer loads but not geometry without strong modelling assumptions; many sensors would be intrusive and expensive, and data are noisy under manoeuvres [62, 63].

Camera-based techniques

- **Digital Image Correlation (DIC):** Full-field strain mapping with dense speckle patterns has shown sub-millimetre accuracy in controlled campaigns on paraglider canopies, but requires extensive surface preparation, short stand-off distances and stable, tunnel-like conditions—impractical for outdoor, full-scale kites [42].
- **Gaussian Splatting (3D-GS):** Useful for fast visualisation of dynamic scenes, yet under-regularised for precise geometry; accuracy degrades with sparse viewpoints. Hybridising with LiDAR/SDF constraints improves depth/mesh fidelity but adds complexity and cost [64, 65].
- **Structured-light 3D scanning:** Delivers sub-millimetre accuracy at short range in controlled, static setups but degrades outdoors on moving, reflective fabric surfaces; range/brightness limits and multi-camera logistics make it unsuitable for a flying kite [66, 67, 68].
- **Depth cameras:** Stereo depth units are easy to deploy but lose precision at moderate–long range due to small baselines; long-range models exist yet remain limited for canopy-scale accuracy and robustness compared with photogrammetry [69, 70].
- **Time-of-Flight (ToF) cameras:** Typical usable ranges are $\sim 10\text{--}15$ m with errors in the order of 1% of distance and strong sunlight sensitivity; high-end systems are costly and still challenged by outdoor dynamics, so they are not competitive for precise kite shape [71, 72, 73].

LiDAR

Vehicle and mapping studies show high temporal coverage and centimetre-scale ranging over tens of metres, but sequential scanning can distort fast motion; weather (rain/fog) reduces returns, and vibration needs mitigation [74, 75, 76, 77, 78]. Overall, LiDAR is feasible for point-cloud envelopes, with residual concerns on occlusion, motion distortion and survivability; costs have decreased but remain non-trivial [79, 76]. Similarly wind-tunnel laser scanners proved accurate on static targets but are unsuited to continuously deforming flight [27].

Ranging techniques

- **Radar:** Compact modules offer long range and very high sampling, yet resolution is typically too coarse for deformation (e.g. 0.56 m CW radar; bistatic errors < 10 m; sub-metre only with wide bandwidth/arrays) [80, 81, 82, 83].
- **Ultrasonic Distance Measuring:** Accurate at short range in static scenes, but air range is usually < 10 m; narrow beams, echo ambiguity, slow wave speed and misalignment make dynamic, metre-scale kite measurements unreliable [84, 85, 86, 87, 88, 89].
- **GPS:** Standard receivers (few-metre errors) are useful for trajectories but not deformations; this holds in AWE field tests with Pixhawk/Xsens integrations [33, 8, 45, 90, 91].
- **RTK GPS:** Carrier-phase corrections yield centimetre-level positions at up to ~ 50 Hz and can reduce infrastructure via NRTK, but still provide sparse nodal geometry rather than full canopy shape and add high costs and logistic difficulties [92, 93, 90].
- **CDGPS:** Shared-antenna, carrier-phase setups enable millimetre-level relative positioning between receivers; promising for precise baselines but hardware topology and multipath limit practicality on a fabric wing [94].
- **Retro-reflective/diffuse/thru-beam photoelectric sensors:** All require tight alignment/line-of-sight and yield only binary or 1D cues; not informative for deforming fabric geometries in flight [62, 63, 95].
- **Draw-wire encoders:** Provide distance but are intrusive, drag-sensitive and bulky for tip-to-tip spans; dynamic, accurate use would be costly with limited shape information [96].

- **Linear potentiometers:** Limited stroke and alignment sensitivity make them unsuitable for large, flexible, moving structures [97].

Structural and kinematic sensing technologies

- **Inertial Measurement Units (IMUs):** Provide orientation/kinematics at sparse nodes and are lightweight, but drift under high dynamics, yaw is sensitive to magnetic disturbance, and rigid mounting/calibration are critical; useful for global trends rather than precise in-plane shape without a strong model [98, 8, 16, 99, 15, 100].
- **Strain tapes:** Resistive/FBG/piezo implementations offer high sampling and local strain with simple integration, yet measure essentially 1D strain, demand careful temperature calibration and durable bonding to fabric, and require dense coverage for reliable shape inference [101, 102, 103, 104, 105, 106, 107].
- **Flex sensors:** Lightweight, low-cost bending sensors that report single-axis curvature; pairs are needed for tension/compression and spatial resolution is low, limiting stand-alone use for complex 3D deformations on flexible skins [108, 109, 110].
- **Fibre Optic Shape Sensors (FOSS):** Distributed curvature sensing (often via FBGs in multi-/multicore fibres) enables millimetre-scale shape over metre lengths at kHz rates and is immune to EMI, but interrogators are costly/heavy and practical sensor length/coverage can be limiting in flight applications [111, 112, 113].

2.6 Summary and outlook

This literature review has examined the aerodynamics of LEI kites, the principal modelling approaches (from quasi-steady aero-structural to higher-fidelity FSI frameworks), and the characteristic deformation modes relevant to performance and control. It has also evaluated prior kite experiments, focusing on state-of-the-art setups while also identifying possible points of improvement. Additionally a broad range of candidate measurement techniques for deformation tracking were investigated, highlighting why accurate outdoor, full-scale shape reconstruction is challenging and which methods remain most promising.

Based on insights from the literature study, discussions with engineers in the field, some knowledge gaps within literature, specifically related to AWES, have been identified:

- (i) **No established guideline for field measurement:** There is no end-to-end, validated protocol for instrumenting flexible, tethered membranes in outdoor flight (at > 300 m altitude and > 35 ms^{-1} apparent wind).
- (ii) **No quantitative in-flight deformation dataset:** Outside controlled facilities, centimetre-level, time-resolved 3D shape data for LEI kites are essentially absent. Reliable measurements of spanwise twist/warp, anhedral, projected-area change, and leading/trailing-edge deflections across powered/depowered states and during steering is lacking. While qualitative data is available, quantitative data is missing.
- (iii) **Model-validation data are scarce:** Coupled aero-structural models need synchronised shape and airspeed measurements on the same flights for calibration and verification; publicly available datasets with this breadth and timing fidelity are rare, and there is no accepted benchmark for parameter identification.

The next chapter sets the research objective and questions, translates the findings into concrete measurement requirements (e.g. accuracy, range, sampling rate, synchronisation, intrusiveness and robustness), and motivates the final selection of techniques to be used in this work.

Conceptual research design 3

This chapter outlines the conceptual design of the study. It first restates the research objective, followed by the research questions in Section 3.1 and Section 3.2 respectively. It then selects the most suitable technique based on the requirements and following trade-off, which are found in Appendix A and Appendix B respectively.

3.1 Research objective

To bridge the gaps identified in the literature review, the following research objective has been formulated:

"To quantify the large scale deformations of a leading-edge inflatable kite used in airborne wind energy through in-flight experimental measurements."

As previously stated in the introduction. Addressing these gaps requires an in-flight measurement approach that is accurate, synchronised, and minimally intrusive in outdoor operation. The present work aims to contribute to research by defining a framework for in flight deformation measurements, providing validation data for FSI models and a better quantitative understanding of kite deformations in AWES.

3.2 Research questions

To achieve this objective, the research addresses both the experimental quantification of kite deformations and their aerodynamic context. The questions are grouped into three themes:

- (i) **Measurement techniques:** Determining the most suitable experimental methods to capture in-flight deformations accurately and reproducibly, with the broader aim of establishing a framework for future deformation studies.
- (ii) **Deformation characterisation:** Understanding how an LEI kite deforms under different flight conditions, in order to form a systematic picture of in-flight deformation behaviour.
- (iii) **Model validation and contribution:** Using experimentally measured deformations to evaluate the predictive capability of existing aero-structural and fluid-structure interaction (FSI) models, and to provide a data foundation for future model development.

Within these themes, the following research questions (RQ) have been defined:

RQ1: What are the most suitable experimental methods for measuring in-flight deformations of LEI kites?

RQ1a: Which techniques are available for measuring in-flight deformations of soft-wing kites?

RQ1b: How do these techniques compare in terms of accuracy, feasibility, and reliability for deformation characterisation?

RQ1c: How can measurements be conducted reproducibly in outdoor environments with inherently variable wind conditions?

RQ1d: How adaptable is the test setup to different kite designs, and would it function across various configurations?

RQ2: How does a LEI kite deform during flight?

RQ2a: What are the primary deformation modes observed in operation?

RQ2b: How do deformations vary under different flight conditions, such as changes in power setting or turning manoeuvres?

RQ2c: How do deformations compare to the CAD model?

RQ3: How can experimentally measured deformations be used to assess and validate aero-structural and FSI models of LEI kites?

RQ3a: Can experimental measurements be used to validate FSI models?

RQ3b: To what extent do FSI models capture real-world deformation behaviour?

3.3 Selection of measurement techniques

To ensure successful execution, a structured set of requirements is defined across the experimental setup, instrumentation, data processing, validation, and practical constraints. Together, these requirements provide a framework for an accurate, reliable, and reproducible measurement system under variable outdoor wind conditions. While formulated for the TU Delft V3 kite (see Chapter 1), they can be adapted to other LEI kites with different geometries or sizes. Not all requirements apply equally to every technique, as some rely on fundamentally different measurement principles. The list of requirements is provided in Appendix A.

Based on these requirements, an extensive trade-off between the measurement techniques found throughout the literature research was performed. The discussion and sensor trade-off is detailed in Appendix B. After this trade-off several measurement techniques remain viable, though their respective performance varies significantly across different criteria.

For **global shape mapping**, stereoscopic photogrammetry is preferred for accurate out-of-plane reconstruction; single-camera photogrammetry is simpler but lacks depth and offers little benefit over stereo. DIC is unsuitable at long range and requires dense patterns. Depth/ToF cameras are short-range and lose precision with distance. LiDAR can capture precise 3D shape but is crash-vulnerable and costly; with protection it remains a contender. Overall, stereo photogrammetry and LiDAR are the most viable.

For **distance and position**, draw-wire sensors perform poorly over longer spans due to sag and drag. CDGPS is strong because it provides relative positions and is robust across environments; for kite-borne nodes, UWB is attractive, with accuracy improvable via filtering. RTK GPS adds little over CDGPS here, so CDGPS and UWB are the preferred options.

Among **structural/kinematic methods**, FOSS would be ideal but is financially out of scope. Strain tapes and flex sensors need dense networks and extensive wiring yet still only provide local measurements, making full-shape inference uncertain. IMUs are the only practical option in this class, but their absolute positioning limits and demonstrated accuracy shortfalls for canopy shape [15] make them a secondary choice.

Based on the trade-off analysis, the most promising techniques (in order) for experimentally measuring the kite's deformation are:

1. **Stereoscopic photogrammetry**
2. **Ultra-wideband modules (UWB)**

As stereoscopic photogrammetry scores the highest in the trade-off, this will be the main technique used in the research. The ultra-wideband modules will be used for ranging as an additional technology. CDGPS was discarded due to the relatively high cost and complexity of the sensors.

Stereoscopic photogrammetry setup 4

This chapter discusses the experimental setup of the stereoscopic photogrammetry pipeline. This includes explanations on the basic principles in Section 4.1, followed by an introduction of the GoPro cameras used to meet this end in Section 4.2. Afterwards, the extensive calibration procedure used is detailed (Section 4.3), which includes both single-camera calibration, stereoscopic calibration and a baseline calculation. With the calibrated system, the choice and performance of markers is analysed in Section 4.4. For these markers, using the calibration matrices, the accuracy of the system is determined in Section 4.5. The chapter ends with the algorithm constructed for automated shape reconstruction (Section 4.6).

4.1 Basic principle

Stereoscopic photogrammetry relies on triangulation and requires at least two cameras. With a single camera, only the two-dimensional (2D) position of a point can be obtained, together with its motion over time. This principle is illustrated in Fig. 4.1. However, its depth relative to the camera cannot be resolved without additional information.

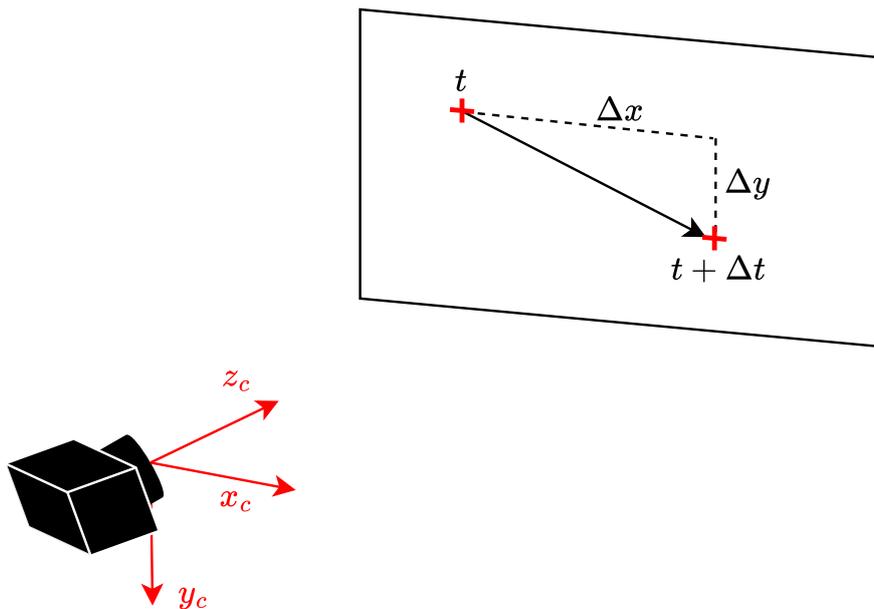


Figure 4.1: Single-camera photogrammetry principle. A point can be tracked in time, but only in 2D. Depth cannot be resolved, as the 2D image plane provides only the x and y coordinates. Determining z requires an additional view of the object.

The main limitation of a single-camera system is the inability to measure distance to an object. In practice, this restricts the method to 2D applications. Some studies, as discussed in the literature review (Chapter 2), have attempted to approximate depth by combining the known size of a marker with the field of view (FOV) of the camera. Nevertheless, such approaches remain significantly less accurate than a stereoscopic setup.

The stereoscopic principle can be compared to human vision. With only one eye open, depth cannot be measured directly: it can only be estimated by relying on prior knowledge of an object's size and by noting that distant objects appear smaller, and vice versa. With both eyes open, the brain infers depth from binocular (retinal) disparity, i.e., the small angular difference between the two retinal

images caused by the eye separation. Similarly, when a second camera is introduced, the relative depth of objects can be determined. Provided that the two images or video frames overlap, depth is resolved by triangulating the 2D positions of the object in both images and by using the known spatial relationship between the cameras. This principle is illustrated in Fig. 4.2 [114].

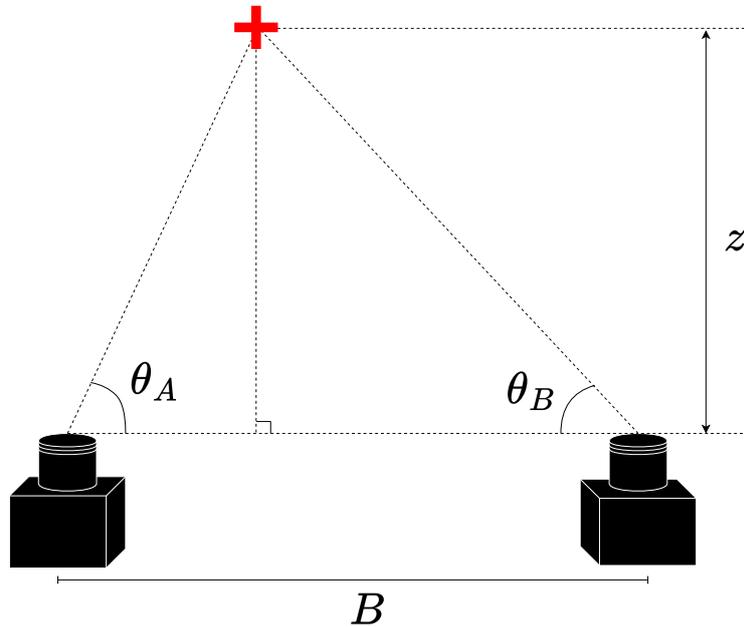


Figure 4.2: Triangulation principle in stereoscopic photogrammetry. After image rectification, discussed later in this chapter, corresponding points in both camera views lie on the same epipolar line. This geometric constraint reduces the problem to two dimensions and enables the 3D coordinates to be solved through simple geometrics.

Once formulated geometrically, the problem reduces to determining the depth of the object from the viewing angles of the two cameras. This is expressed mathematically as:

$$B = \frac{z}{\tan(\theta_A)} + \frac{z}{\tan(\theta_B)} = z \left(\frac{\cos(\theta_A)}{\sin(\theta_B)} + \frac{\cos(\theta_B)}{\sin(\theta_A)} \right) = z \frac{\sin(\theta_A + \theta_B)}{\sin(\theta_A)\sin(\theta_B)},$$

hence:

$$z = B \frac{\sin(\theta_A)\sin(\theta_B)}{\sin(\theta_A + \theta_B)},$$

where B denotes the (fixed) baseline between the cameras. This formulation provides the depth z of the object, which, when combined with its in-plane motion, yields the full three-dimensional position of the point in space. In practice, these geometric principles are implemented computationally using the OpenCV library [115] in Python:

```
1 import cv2
```

OpenCV is an open-source computer vision toolkit that provides efficient algorithms for camera calibration, stereo rectification, correspondence matching, and triangulation. It enables the theoretical framework described above to be translated into a practical pipeline for extracting 3D coordinates from stereo image data.

4.2 GoPro Hero9 Black

For this study, two GoPro Hero9 Black cameras¹ were selected as the stereoscopic imaging system. The choice is motivated by their light weight, robustness, and ready availability. The camera is shown in Fig. 4.3.

¹<https://gopro.com/en/us/>



Figure 4.3: GoPro Hero9 Black from three different angles. Two such cameras are used in the stereoscopic photogrammetry setup [116].

A potential advantage of this model is the integrated stabilisation software, designed to suppress vibrations and maintain horizontal levelling. However, these functions reduce the effective number of pixels in the field of view and can cause inconsistencies in pixel location between two devices. As stereoscopic photogrammetry relies on precise correspondence between images, such functions had to be disabled for the present study.

The wide-angle lens variant of the Hero9 Black was used, as it offers the same video quality and frame rate as the linear lens, while increasing the likelihood that the entire kite remains within the field of view during testing. The most relevant specifications of the camera are summarised in Table 4.1.

Table 4.1: Relevant parameters of the GoPro Hero9 Black. Only specifications relevant to this study are listed, considering the highest video quality setting to ensure maximum accuracy [117].

Parameter	Value
Video quality	5K
Horizontal resolution	5120 pixels
Vertical resolution	2880 pixels
Frame rate	30 fps
Maximum field of view	$118 \times 69^\circ$
Mass	158 g

The field of view of the camera satisfies requirements REQ-INS-06b and REQ-INS-06c, confirming the suitability of the equipment for global shape mapping technologies, as specified in REQ-INS-06. In addition, the sampling frequency of 30 frames per second is more than sufficient for the dynamics of this study, fulfilling REQ-INS-04.

Finally, the Hero9 Black is designed for high-intensity sports, where crashes and harsh environmental conditions are expected. Its construction therefore provides resilience against impacts and severe weather. For additional protection, both cameras were mounted in reinforced housings, addressing requirements REQ-INS-05, REQ-INS-10, and REQ-INS-11.

4.3 Camera calibration

Calibration is required to determine the intrinsic and extrinsic parameters of the camera system. This process establishes the mapping between 3D real-world coordinates (X, Y, Z) and their 2D pixel projections on the image plane (x, y, z) , where z is always equal to one, resulting in $(x, y, 1)$.

The intrinsic parameters are unique to each camera and lens. They include the focal lengths (f_x, f_y) , the optical centre (c_x, c_y) , and the distortion coefficients (radial and tangential). These parameters define how the camera internally forms an image and are obtained through single-camera calibration (see Section 4.3.1).

The extrinsic parameters describe the spatial relationship between the cameras, i.e. their relative orientation and translation in 3D space. They are not camera-specific but define how the two views relate to one another in a stereo system. For individual calibration frames, the extrinsics encode the rotation and translation between the calibration target and the camera. In stereoscopic calibration, however, the extrinsic parameters define the fixed transformation between the two cameras, and thus remain constant once the setup is rigidly mounted. These are obtained through the stereo calibration process, as discussed in Section 4.3.2.

4.3.1 Single-camera calibration

Intrinsic calibration is necessary because the pinhole camera model assumes an ideal projection in which all light rays pass through a single point and are projected linearly onto the image plane. Real camera lenses deviate from this model due to manufacturing tolerances, wide-angle optics, and lens distortions. Without intrinsic calibration, image measurements cannot be accurately related to real-world geometry, which would compromise triangulation accuracy. The main components of intrinsic calibration are illustrated in Fig. 4.4.

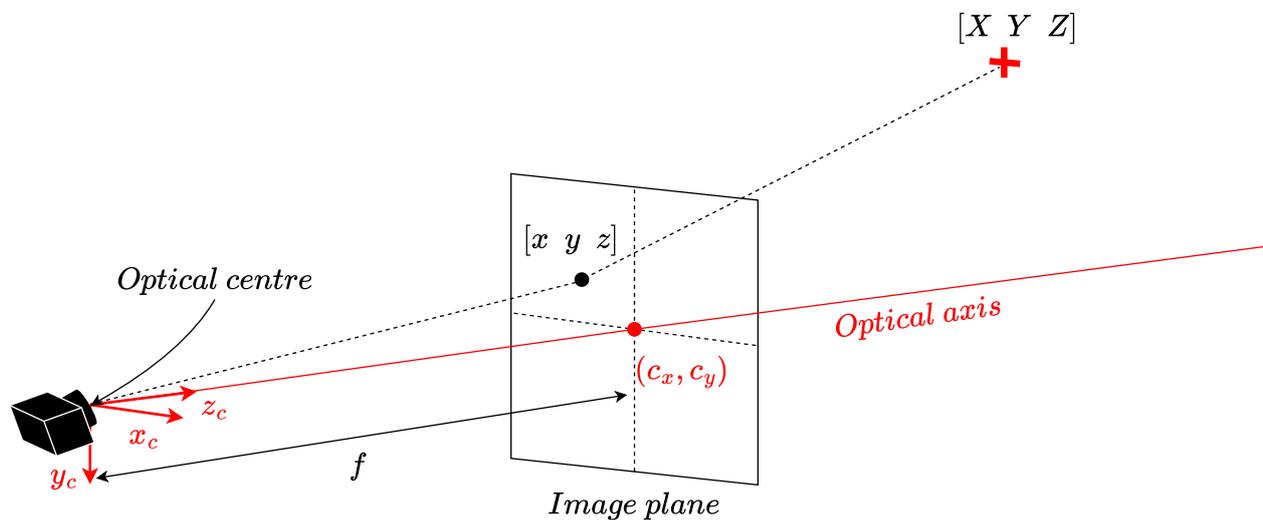


Figure 4.4: Intrinsic calibration parameters estimated in OpenCV: focal lengths (f_x, f_y) , principal point (c_x, c_y) , and distortion coefficients describing deviations from the ideal pinhole projection.

Focal length and optical centre

The focal length determines how strongly the camera maps real-world points to the image plane. In digital cameras, f_x and f_y are expressed in pixels, effectively scaling world coordinates into the image frame. The optical centre (c_x, c_y) , also called the principal point, is the location in the image plane where the optical axis intersects. This results in zero distortion along the optical axis. The principal point is usually near the geometric centre of the image, though small offsets occur in practice. For 5K video resolution $(5120 \times 2880 \text{ px})$, the optical centre is therefore expected around $(2610, 1440) \text{ px}$.

This mapping from real-world coordinates (X, Y, Z) to pixel coordinates (x, y) is expressed through the camera matrix \mathbf{K} :

$$Z \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \underbrace{\begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{K}} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix},$$

where \mathbf{K} is the intrinsic camera matrix. This formulation ensures that 3D coordinates in the camera frame are projected consistently onto the 2D image plane. In reality, reconstruction (i.e. recovering

3D coordinates from 2D image points) requires additional information, such as depth estimation or triangulation in a stereo setup.

Lens distortion

Real lenses introduce distortions that must be compensated for in order to achieve accurate geometric reconstruction. These are typically decomposed into radial and tangential distortion.

Radial distortion occurs because light rays passing through the periphery of the lens bend more strongly than those near the optical axis. This causes effects such as barrel distortion (straight lines bulging outwards) or pincushion distortion (lines bending inward). OpenCV models radial distortion as:

$$\begin{aligned}x_{r,\text{dis}} &= x \left(1 + k_1 r^2 + k_2 r^4 + k_3 r^6 \right), \\y_{r,\text{dis}} &= y \left(1 + k_1 r^2 + k_2 r^4 + k_3 r^6 \right),\end{aligned}$$

where r is the radial distance from the optical centre, in normalised image coordinates. The coefficient k_1 dominates, while k_2 and k_3 provide higher-order corrections [118]. For wide-angle lenses, such as the GoPro Hero9 Black, k_1 is typically negative and large in magnitude (i.e. $k_1 < -0.2$) [119]. For linear lenses, the k_i values approach zero.

Tangential distortion arises from slight misalignments of the lens with respect to the image sensor, and is modelled as:

$$\begin{aligned}x_{t,\text{dis}} &= x + \left(2p_1 xy + p_2 (r^2 + 2x^2) \right), \\y_{t,\text{dis}} &= y + \left(p_1 (r^2 + 2y^2) + 2p_2 xy \right),\end{aligned}$$

where (x, y) are the ideal normalised coordinates and $(x_{\text{dis}}, y_{\text{dis}})$ are their distorted counterparts [118]. Since modern manufacturing aligns lenses precisely, the tangential distortion coefficients p_1 and p_2 are typically close to zero. Combining radial and tangential effects yields the complete OpenCV distortion model:

$$\begin{aligned}x_{\text{dis}} &= x \underbrace{\left(1 + k_1 r^2 + k_2 r^4 + k_3 r^6 \right)}_{\text{radial}} + \underbrace{\left(2p_1 xy + p_2 (r^2 + 2x^2) \right)}_{\text{tangential}}, \\y_{\text{dis}} &= y \underbrace{\left(1 + k_1 r^2 + k_2 r^4 + k_3 r^6 \right)}_{\text{radial}} + \underbrace{\left(p_1 (r^2 + 2y^2) + 2p_2 xy \right)}_{\text{tangential}}.\end{aligned}$$

OpenCV represents the distortion parameters as a 1×5 vector:

$$\mathbf{d} = (k_1 \quad k_2 \quad p_1 \quad p_2 \quad k_3).$$

For the GoPro Hero9 Black, the distortion coefficients of the wide lens, used in this research, are expected to be substantially larger in absolute value compared to a linear lens, reflecting the increased radial distortion.

Calibration procedure

To obtain reliable estimates of the intrinsic camera parameters, images of known calibration patterns must be taken from multiple viewpoints. OpenCV provides built-in routines for this purpose [118], supporting a range of pattern types. In this study, three different calibration boards were considered: a chequerboard, a ChArUco board, and an asymmetric circle grid, which are shown in Fig. 4.5.

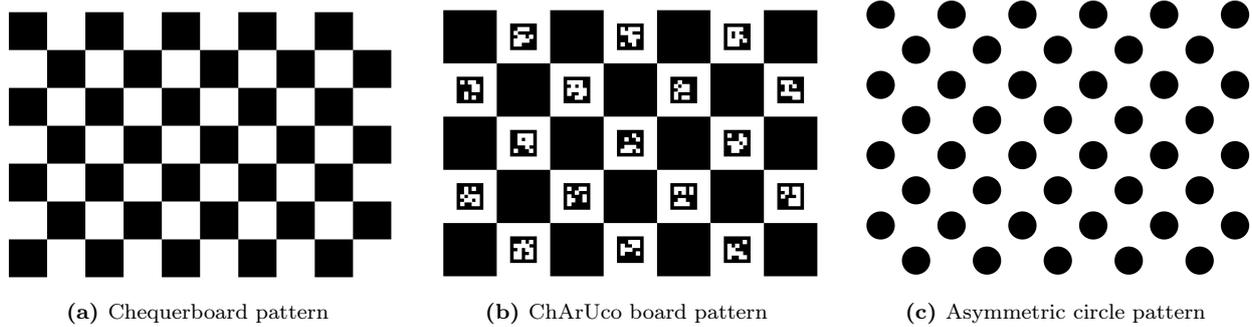


Figure 4.5: Calibration patterns used in this study [118]. These patterns were printed on A3 paper to maximise visibility.

These patterns are suitable for this purpose because their geometry is precisely known: both the layout (e.g. number of rows and columns) and the dimensions (e.g. square size, marker size, or circle diameter) are fixed. The pattern specifications are summarised in Table 4.2. By detecting the feature points (corners or circle centres) in an image, their corresponding 3D locations on the calibration board are also known. This one-to-one mapping between image points and real-world points allows the intrinsic parameters of the camera to be estimated through optimisation.

Table 4.2: Specifications of the calibration patterns used.

Pattern	Grid size	Dimensions	Number of detectable features
Chequerboard	9×6	39.0×39.0 mm per square	54
ChArUco board	7×5	55.5×55.5 mm per square 28.0×28.0 mm per marker	41
Asymmetric circles	11×4	25.0 mm diameter per circle	44

For the calibration procedure, a total of 19 frames were collected per camera using the 5K wide-angle setting of the GoPro Hero9 Black. Of these, 11 frames were taken with the calibration board placed perpendicular to the optical axis, ensuring that the board covered different regions of the image plane, including all four corners. An additional 8 frames were captured at an angle, with views taken from each corner and side, in order to incorporate strong perspective effects. Capturing angled views is particularly important to improve the estimation of radial distortion coefficients, as distortion effects are most pronounced near the periphery of the lens. In practice, although a chequerboard pattern can theoretically be calibrated with as few as two images, more images with varied viewpoints are recommended to achieve robust and noise-resilient parameter estimation.

OpenCV provides dedicated functions to detect and process the calibration patterns:

```

1 #For chequerboard calibration:
2 cv2.findChessboardCorners()
3 cv2.calibrateCamera()
4
5 #For ChArUco board calibration:
6 cv2.aruco.detectMarkers()
7 cv2.aruco.interpolateCornersCharuco()
8 cv2.aruco.calibrateCameraCharuco()
9
10 #For asymmetric circle grid calibration:
11 cv2.findCirclesGrid()
12 cv2.calibrateCamera()

```

An example of the detected corners using the chequerboard pattern is shown in Fig. 4.6a. The barrel distortion of the GoPro wide lens is clearly visible in the input image.

Single-camera calibration results

The accuracy of the intrinsic calibration can be quantified using the reprojection error [120]. This metric expresses how closely the image points predicted by the calibrated model match the points

actually detected in the calibration images. After calibration, the known 3D object points \mathbf{X}_i from the calibration board are projected back onto the image plane using the estimated intrinsic matrix \mathbf{K} and distortion coefficients \mathbf{d} . The resulting 2D projections $\hat{\mathbf{x}}_i$ are compared to the detected image points \mathbf{x}_i . The reprojection error ϵ_r is defined as the average Euclidean distance between the detected and reprojected points, in pixels:

$$\epsilon_r = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|,$$

where N is the total number of detected image points. A lower value of ϵ_r indicates a more accurate intrinsic calibration. This error was computed in OpenCV using:

```
1 cv2.projectPoints()
```

averaged across all calibration images. As a qualitative validation of the distortion model, Fig. 4.6b shows the same checkerboard frame undistorted using the estimated intrinsic matrix and distortion coefficients; note the straightened grid and cropping to a rectangular aspect.

The intrinsic calibration results are summarised in Table 4.3 and Table 4.4. In practice, the checkerboard and ChArUco board patterns produced reprojection errors well below 0.2 px, indicating excellent calibration accuracy in line with expectations [121]. The asymmetric circle grid proved less reliable: not all frames were successfully detected, particularly those with the board at the image periphery, and only three usable frames were recognised for the second camera. This limited dataset resulted in high reprojection errors and inconsistent intrinsic parameters, making the circle grid calibration results unsuitable for the present setup.

In both cameras the checkerboard pattern consistently yielded the lowest reprojection errors, confirming it as the most robust method for calibration. Furthermore, the wide-angle lens exhibited the expected strong negative radial distortion coefficient k_1 , consistent with theoretical behaviour of wide-angle optics. While the parameters for the two cameras are similar, slight differences remain, which can be attributed to manufacturing tolerances and internal processing. Separate calibration of each camera is therefore essential to ensure accuracy.

Based on these results, the intrinsic parameters obtained using the checkerboard pattern were selected for both cameras and are used in the subsequent analysis. The final intrinsic matrices \mathbf{K} and distortion vectors \mathbf{d} are given in Table 4.5. Interestingly, no tangential distortion occurs ($p_1 = p_2 = 0$), meaning the lenses are perfectly aligned with the image plane, as expected.

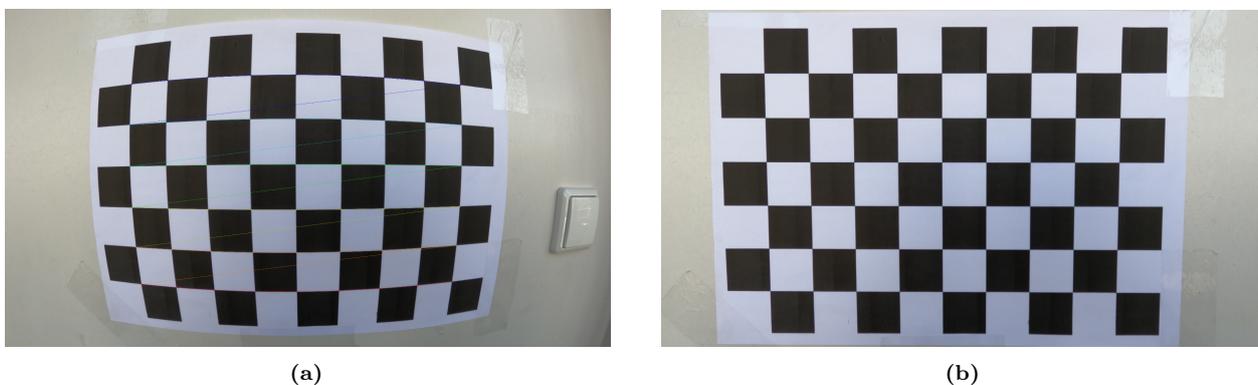


Figure 4.6: Effect of lens distortion correction on a checkerboard calibration image. (a) Corner detection in OpenCV on the checkerboard pattern (GoPro Hero9 Black, 5K wide-angle). Barrel distortion is visible near the image edges. (b) Undistorted checkerboard using the estimated intrinsics and distortion coefficients (same frame as (a)). Barrel distortion is removed; the result is cropped to a rectangular aspect.

Table 4.3: Calibration results for camera 1. The reprojection error confirms that calibration with the chequerboard pattern was the most accurate.

Camera 1			
Calibration method	Chequerboard	ChArUco board	Asymmetric circles
f_x	2340.749	2353.379	1954.809
f_y	2347.773	2361.657	1702.237
c_x	2555.872	2561.512	1391.997
c_y	1466.105	1460.423	1306.622
k_1	-0.252	-0.264	-0.188
k_2	0.080	0.097	0.048
p_1	0	0	0
p_2	0	0	-0.019
k_3	-0.013	-0.019	-0.006
Reprojection error ϵ_r (px)	0.153	0.199	2.689

Table 4.4: Calibration results for camera 2. The reprojection error again shows that the chequerboard pattern produced the most reliable results.

Camera 2			
Calibration method	Chequerboard	ChArUco board	Asymmetric circles
f_x	2337.803	2337.097	18681.800
f_y	2344.095	2343.749	14576.767
c_x	2568.016	2570.036	2697.077
c_y	1482.403	1468.156	1493.556
k_1	-0.253	-0.264	-10.802
k_2	0.083	0.097	496.685
p_1	0	0	0.190
p_2	0	0	-0.031
k_3	-0.014	-0.019	-11176.886
Reprojection error ϵ_r (px)	0.159	0.186	5.398

Table 4.5: Final intrinsic parameters for both cameras

Variable	Camera 1	Camera 2
f_x	2340.749	2337.803
f_y	2347.773	2344.095
c_x	2555.872	2568.016
c_y	1466.105	1482.403
k_1	-0.252	-0.253
k_2	0.080	0.083
p_1	0	0
p_2	0	0
k_3	-0.013	-0.014

4.3.2 Stereoscopic calibration

In addition to intrinsic calibration per camera, stereoscopic photogrammetry requires a stereo (extrinsic) calibration to recover the relative pose between the two cameras. Stereo calibration does not re-estimate intrinsics; instead, it estimates the rigid transform from the left to the right camera. Although the cameras are mounted with a nominal baseline B and approximately parallel optical axes, small mounting tolerances and the unknown location of each sensor within its housing introduce pose uncertainty. Stereo calibration is therefore required to obtain an accurate relative rotation and translation, which in turn enables metrically valid depth estimation from disparities.

Baseline calculation

The baseline B —i.e. the centre-to-centre camera spacing—must be chosen to balance depth precision against the common field-of-view. Increasing B enlarges the angle of triangulation, which in turn increases the horizontal disparity observed between the left and right images. Horizontal disparity is defined as the horizontal difference between the projections of the same 3D point in the two image planes (in pixels). As shown in Fig. 4.7, points closer to the cameras project with a larger horizontal disparity than points positioned further away ($z_1 < z_2, \delta_{x,2} < \delta_{x,1}$).

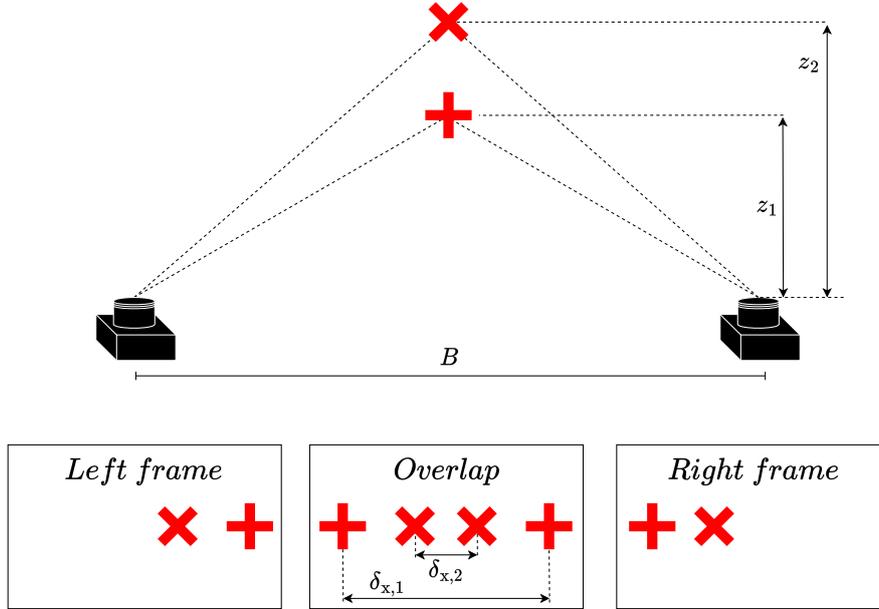


Figure 4.7: Baseline B and disparity concept. For a fixed setup, the closer point at distance z_1 yields a larger horizontal disparity $\delta_{x,1}$ than the further point at z_2 with horizontal disparity $\delta_{x,2}$. Increasing the baseline magnifies horizontal disparities for all depths, thereby improving depth resolution.

A larger baseline therefore improves depth accuracy, since even small differences in depth result in measurable disparity shifts. However, as B increases, the overlap between the two camera views decreases, and parts of the wing may fall outside one camera’s field of view. In addition, a physically wider rig becomes more intrusive and less practical. Thus, the baseline must be selected as a compromise: large enough to provide sufficient disparity for accurate triangulation, yet sufficiently small to maintain full coverage of the wing and to keep the measurement system manageable.

This trade-off can be quantified with the standard first-order depth-error relation [122]:

$$\epsilon_z = \frac{z^2}{Bf} \epsilon_{\delta,x}, \quad (4.1)$$

where z is depth (m), B the baseline (m), f the focal length (pixels), and $\epsilon_{\delta,x}$ the disparity error (pixels). Typical $\epsilon_{\delta,x}$ is ~ 0.25 – 1 px with sub-pixel matching (e.g. SGBM) [122, 123]; extreme cases of poor calibration can increase this (reports of 2 px exist [124]), but filtering and careful calibration make such cases unlikely.

Because $f_x \neq f_y$, there is no unique scalar f in pixels. To be conservative the smaller of (f_x, f_y) is considered, maximising ϵ_z , as $\epsilon_z \propto 1/f$. The focal length used is provided in 2337.803 px.

Using Eq. (4.1) with $z = 7$ m and a range of $\epsilon_{\delta,x}$, the estimated depth error over B is shown in Fig. 4.8. The requirement (REQ-EXP-01b) set the depth accuracy to 5 cm. The minimum baseline that satisfies this requirement in the worst-case assumed $\epsilon_{\delta,x}$ is indicated.

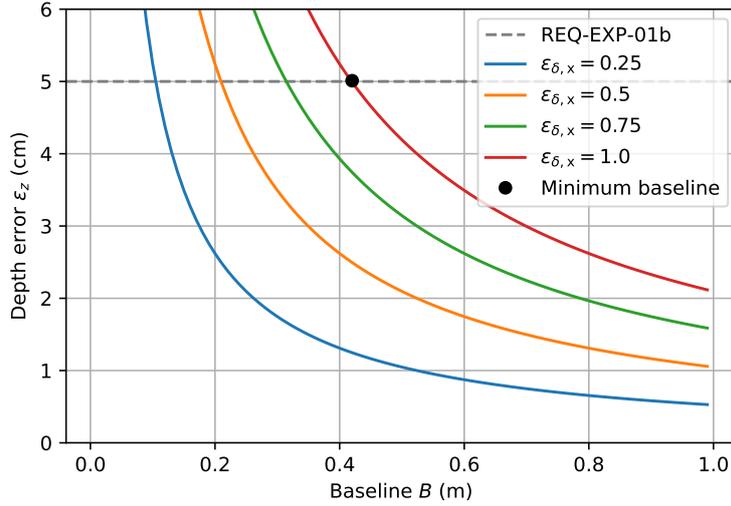


Figure 4.8: Estimated depth error ϵ_z at $z = 7$ m vs. baseline B for several disparity errors $\epsilon_{\delta,x}$. The marked point indicates the design minimum baseline that just meets REQ-EXP-01b in the worst case.

To incorporate additional uncertainty, a safety factor (SF) of 2 was applied to the minimum baseline, provided the kite remains fully within both fields of view. The adopted baseline is listed in Table 4.6.

Table 4.6: Baseline for the V3 test cases at $z = 7$ m. A value of $SF = 2$ was applied to the computed minimum.

z (m)	B (cm)	$SF \cdot B$ (cm)
7	42	84

Rotation and translation matrices

For stereo calibration, image pairs of a fixed calibration target (here, a chequerboard) are captured simultaneously by the left and right cameras. Given the known intrinsic parameters, corresponding points in both image planes constrain the relative pose (\mathbf{R}, \mathbf{t}) via epipolar geometry. To do this, 50 images were taken of the chequerboard pattern with the cameras fixed on a bar with distance $B = 84$ cm. These images were taken at different angles and distances, ensuring complete and accurate extrinsic parameters. Frames in which not all corners in the matching images were detected, were removed. Because intrinsic reprojection errors were low, the intrinsic parameters were held fixed during stereo calibration, and only the extrinsic parameters were computed. The estimated relative rotation and translation take the form

$$\mathbf{R} = \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix}, \quad \mathbf{t} = \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix},$$

with t_x expected to be close to the physical baseline B and \mathbf{R} close to the identity matrix \mathbf{I} for a nearly parallel mounting.

Combining intrinsics \mathbf{K} and extrinsics $[\mathbf{R} | \mathbf{t}]$, the mapping from a 3D point $\mathbf{X} = [X, Y, Z, 1]^\top$ to its homogeneous image point $\mathbf{x} = [x, y, 1]^\top$ is

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \mathbf{K} [\mathbf{R} | \mathbf{t}] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_{11} & r_{12} & r_{13} & t_x \\ r_{21} & r_{22} & r_{23} & t_y \\ r_{31} & r_{32} & r_{33} & t_z \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}.$$

Here x, y denote undistorted image coordinates. Given corresponding points (x_L, y_L) and (x_R, y_R) and the calibrated models, triangulation yields the 3D point; depth (the Z coordinate) follows from the disparity. In practice OpenCV's stereo pipeline is used to compute \mathbf{R} and \mathbf{t} :

```
1 cv2.stereoCalibrate()
```

As an illustration of the stereo-calibration input, Fig. 4.9 shows simultaneously captured left–right views of the checkerboard used to estimate \mathbf{R} and \mathbf{t} .

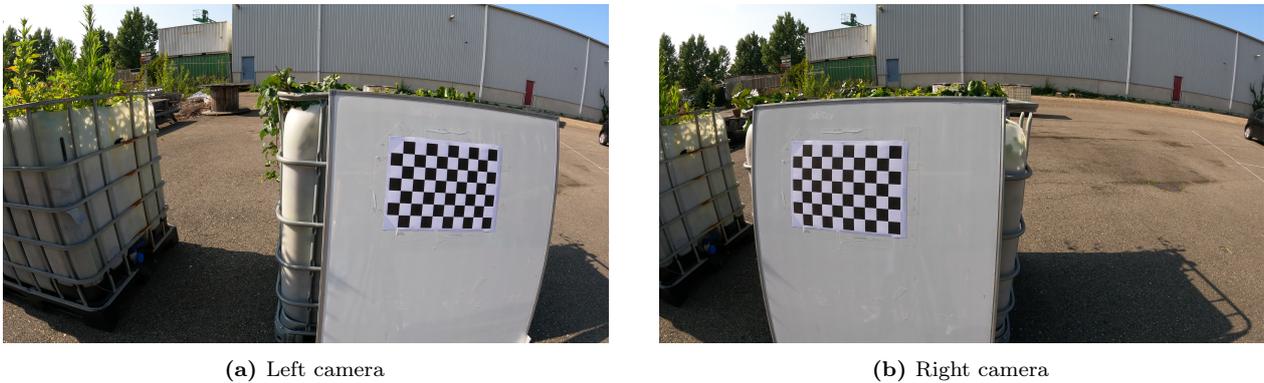


Figure 4.9: Example of the stereoscopic calibration input images (checkerboard). This is one of the pairs used in OpenCV’s `stereoCalibrate` to estimate the relative rotation \mathbf{R} and translation \mathbf{t} between the cameras.

Rectification

Rectification is the process of transforming the stereo image pair so that corresponding epipolar lines, which are the lines connecting corresponding object points in both frames, become perfectly horizontal and aligned. In an unrectified stereo setup, a 3D point projects to two different image points, and its possible matches in the other image must lie somewhere along a slanted epipolar line. Searching along these arbitrary lines is computationally expensive and prone to error. By rectifying the images, the epipolar lines are rotated so they become horizontal and collinear in both images. This means that corresponding points now lie on the same image row, reducing the search for matches to a one-dimensional problem and greatly improving robustness. This thus requires application of both intrinsic and extrinsic parameters. Prior to rectification, the images are undistorted using distortion vector \mathbf{b} . These steps are provided in Fig. 4.10.

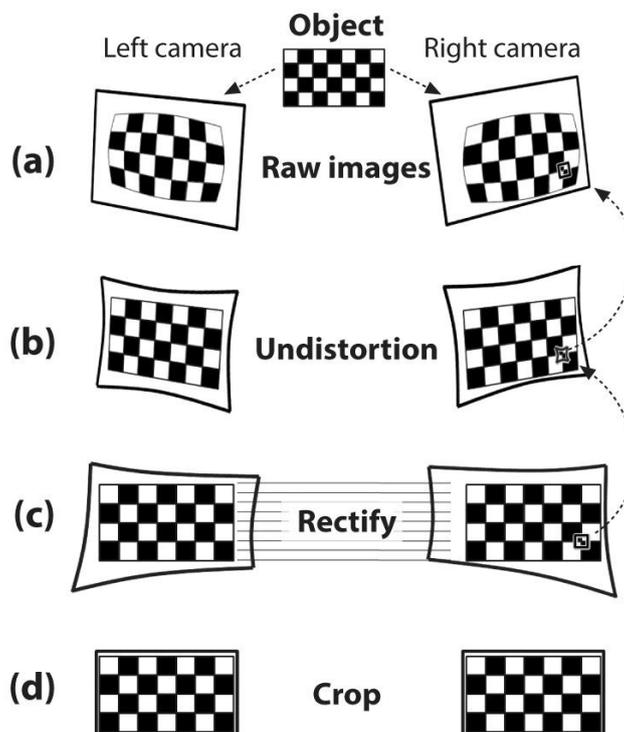


Figure 4.10: Rectification pipeline, adapted from [125]. (a) Raw stereo images; (b) undistortion removes lens distortion; (c) rectification aligns epipolar lines horizontally; (d) the images are cropped to retain rectangular frames.

Stereoscopic calibration results

The stereo calibration provides the relative rotation and translation between the two cameras, which together define the baseline. These extrinsic parameters are reported in Table 4.7. As expected, the rotation is close to identity and the x -translation magnitude approximates the physical baseline ($B = 84$ cm), supporting the validity of the extrinsic solution.

Table 4.7: Stereo extrinsics for the adopted setup ($B = 84$ cm), rounded to three decimals.

Extrinsic parameters	
r_{11} (–)	1.000
r_{12} (–)	0.007
r_{13} (–)	0.003
r_{21} (–)	-0.007
r_{22} (–)	1.000
r_{23} (–)	-0.002
r_{31} (–)	-0.003
r_{32} (–)	0.003
r_{33} (–)	1.000
t_x (m)	-0.835
t_y (m)	-0.007
t_z (m)	0.004

To quantify the quality of the stereo calibration, the epipolar error is the most relevant diagnostic. Rectification aligns epipolar lines so that corresponding points lie on the same horizontal image row (Fig. 4.10). The epipolar error (ϵ_e) is computed by projecting detected corner points from one image onto their corresponding epipolar lines in the other image and measuring the average distance in pixels. A small error means that the rectification correctly enforces the epipolar constraint, which is essential for triangulation and stereo matching. Unlike frame-wise estimates of rotation or translation differences, which are unstable when derived from a single checkerboard pair, the epipolar error directly quantifies how well the rectified geometry supports point matching. It therefore provides a meaningful and practical measure of whether the calibration is accurate for photogrammetry applications.

Table 4.8 reports the average and maximum epipolar error across all calibration images. The average error of 0.318 px and maximum below one pixel confirm that the calibration is of high quality and that the rectified images are suitable for reliable depth estimation [122, 123].

Table 4.8: Stereo calibration quality metric (averaged across calibration images).

Metric	Value
Average epipolar error $\epsilon_{e,\text{avg}}$ (px)	0.346
Maximum epipolar error $\epsilon_{e,\text{max}}$ (px)	0.962

4.4 Markers

Tracking distinct points on the kite surface is essential for stereoscopic photogrammetry, as the wing itself is largely uniform in colour and lacks natural features. Artificial markers provide robust features that can be detected, tracked, and matched reliably across frames. This section will elaborate on the marker setup used in this experiment.

4.4.1 Size and type

Several studies employ ArUco markers, which not only provide a unique ID but also allow for pose estimation using OpenCV. The drawback, however, is that ArUco codes require a relatively large number of pixels to be detected, due to their internal binary patterns. This makes them impractical for very small surfaces, such as the struts of the kite.

To explore alternatives, simpler red geometric shapes were tested alongside ArUco codes: circles and crosses (see Fig. 4.11). The red colour was chosen deliberately to avoid false positives, as the kite canopy already contains black and white areas. Red stands out strongly against both and is therefore robust to misdetection. The hypothesis is that circles may provide good detection even at small sizes, while crosses might prove more robust to tilt or partial occlusion.

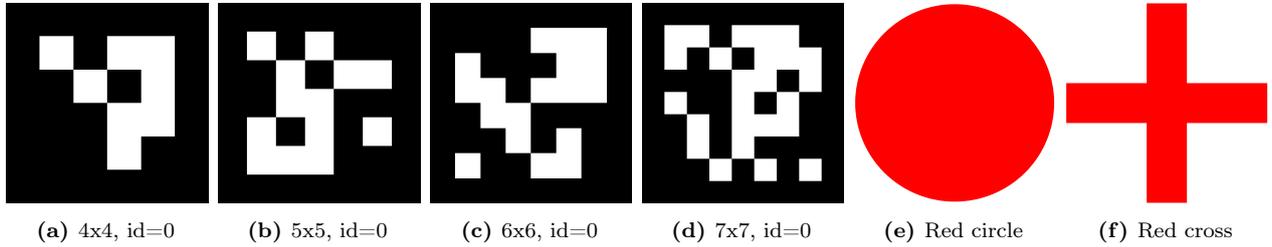


Figure 4.11: Markers tested for detection analysis. The minimum pixel size for detection of each of these markers was determined through experiments. This pixel size could then be translated to a required marker size at distance z .

To quantify minimum detection requirements, videos were recorded with markers placed both centrally and near the edges of the frame. The camera-marker distance was increased stepwise, and the pixel size of the detected marker was logged until detection failed. The largest minimum pixel size across all frame positions was taken as the threshold for reliable detection. With the known wide lens field-of-view (FOV) and image resolution, this pixel requirement is converted into a physical marker size at the single design distance $z = 7$ m (see Fig. 4.12).

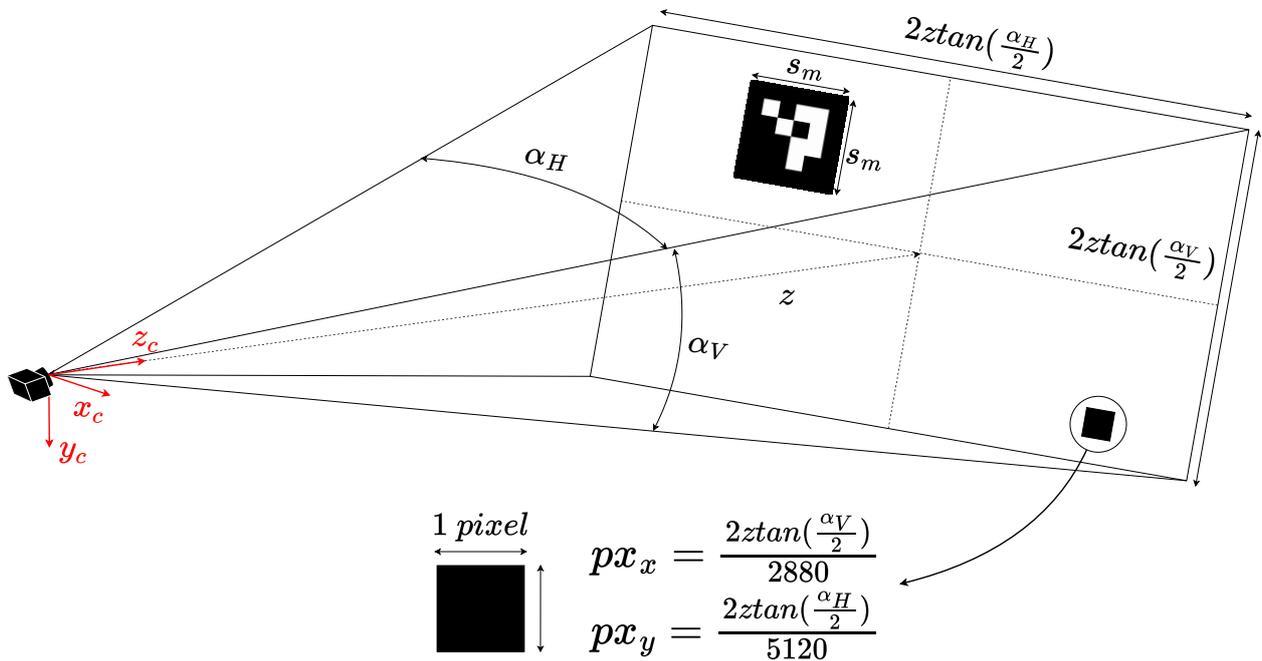


Figure 4.12: Pixel scaling at a given distance z . α_H and α_V denote the horizontal and vertical FOV of the wide lens. The physical size represented by one pixel is px_x and px_y . For an ArUco, s_m is the side length; for a circle, s_m is the diameter. To be conservative, the conversion to millimetres uses $\max(px_x, px_y)$.

Although pixels are square in the image plane, they may correspond to slightly rectangular patches in the scene depending on FOV asymmetry, and are determined using:

$$px_x = \frac{2z \tan(\alpha_H/2)}{W}, \quad px_y = \frac{2z \tan(\alpha_V/2)}{H}.$$

Where px_x and px_y are the sizes in pixels, z the depth from the camera, and α_H and α_V form the FOV. With the wide lens, the horizontal FOV is larger than the vertical; therefore we conservatively use the larger of (px_x, px_y) when converting from pixels to millimetres.

The minimum marker sizes s_m are summarised in Table 4.9. Some variation between ArUco families is expected in practice due to thresholding, lighting, and detector tuning; the finding that a 6×6 code triggered at slightly smaller sizes than a 4×4 is attributed to such run-to-run variability. The red circle is detectable at very small pixel sizes, while the red cross needs slightly more pixels but proved more robust under tilt or partial deformation.

Table 4.9: Minimum detectable sizes for each marker type. Millimetre sizes are computed at the design distance $z = 7$ m using Fig. 4.12.

Metric	4x4	5x5	6x6	7x7	Red circle	Red cross
Min detectable size (px)	27	28	21	34	2	8
Size s_m at $z = 7$ m (mm)	100	100	73	132	9	36

These results indicate that full-size ArUco markers are impractically large for mounting on small kite components such as struts. For smaller surfaces, simpler shapes (red circles or crosses) are preferable, as they remain reliably detectable at much smaller physical sizes.

Sub-pixel accuracy

Large markers allow for sub-pixel accuracy in centre detection, as shown in Fig. 4.13. While each pixel has a single colour, pixels at the boundary of a marker will display a blend of marker and background colour, from which the edge position can be estimated with sub-pixel precision. However, this only works if the marker is entirely in view, which also depends on lighting, saturation and image quality. This improves the accuracy of the detected centre, ensuring positional consistency between the left and right frame.

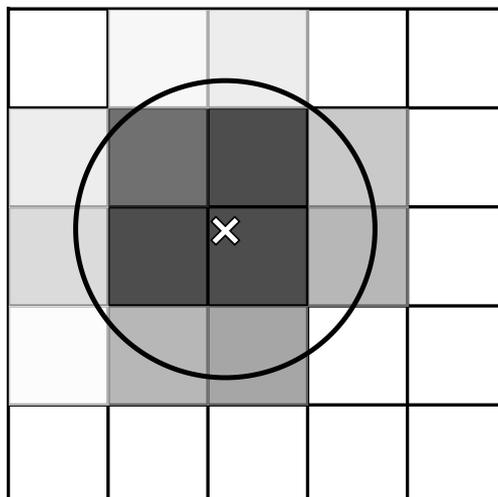


Figure 4.13: Principle of sub-pixel accuracy; circular markers used as example. Boundary pixels contain a blend of marker and background colour, allowing the true marker centre to be estimated more precisely than one pixel.

Crosses require a slightly larger size for detection than circles, they are more robust under non-ideal conditions such as bending, twisting, or oblique viewing angles. For a deformable wing such robustness is particularly important. This, together with the manufacturability, is the reason crosses are used in this research. These markers, presented both on white and black backgrounds, are visualized in Fig. 4.14. The crosses consist of two perpendicular lines of length $s_m = 50$ mm and thickness $t_m = 12.5$ mm

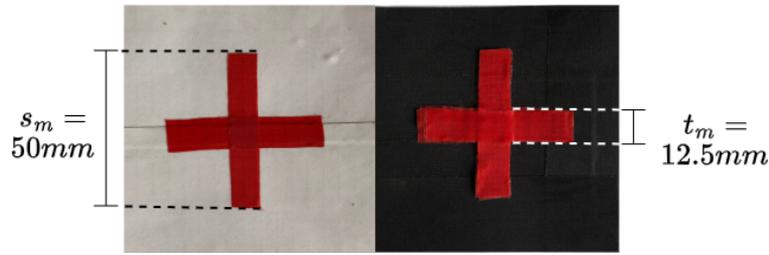


Figure 4.14: Red cross markers on both white (left) and black (right) background.

4.4.2 Placement

In the present V3 configuration, markers are placed such that these enable full 3D reconstruction of the inflatable wing shape rather than only sampling bridle–wing interaction points, which are commonly modelled nodes in FSI models. This denser layout improves observability of deformations across the LE and struts, yielding a more complete and accurate dataset for stereoscopic triangulation and subsequent surface fitting. Additionally, missed detections are less concerning. This redundancy also leads to less uncertainty in the strut and LE deformation, which are the basis for angular extrapolations made in the results.

With the cameras mounted in the bridle system (in the power lines, pointing upward, see Chapter 7), only the underside of the wing is visible—which is also where the bridle–wing interaction points are located.

As illustrated in Fig. 4.15, the LE carries 24 markers with increased density toward the tips to capture the stronger spatial gradients near the wing ends. Table 4.10 shows the number of markers on each of the struts. The markers are made from Dacron repair tape so they remain adhered to the wing through deflation/inflation cycles, improving test-to-test repeatability.

No markers are placed on the canopy of the wing. The canopy shape is much more complex and would therefore require a high spatial density of markers, increasing the manual labour and risk of mismatches significantly. The shape of the canopy can simply be modelled based on the position and orientation of the surrounding struts, and was therefore deemed unnecessary. The current distribution thus provides a complete and well-conditioned dataset that facilitates precise reconstruction of the instantaneous V3 kite shape.

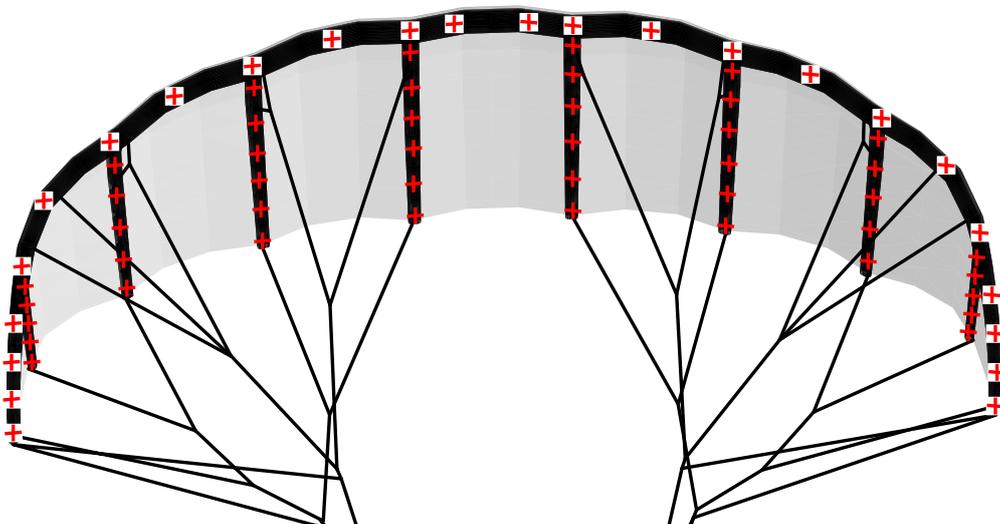


Figure 4.15: Markers used and their placement on the V3. Markers on the struts have a black background, whereas markers on the LE have a white background.

Table 4.10: Marker counts per strut on the V3 (measured from the centreline outward).

Strut (from centre)	# markers
1	6
2	6
3	5
4	5

4.5 Accuracy analysis

To quantify the accuracy of the stereoscopic calibration—and, by extension, the accuracy of the 3D coordinates used in the experiments—an accuracy test is required (REQ-DPV-01). Designing a reliable test is non-trivial in stereoscopic photogrammetry. A particular challenge is that a LEI kite is a non-rigid body, so the exact locations of points on the wing vary and are not known with high precision, making accuracy determination in post-processing difficult.

In a stereoscopic setup, the coordinate frame is typically defined from the image plane of one of the cameras. The z -axis is normal to this plane, while x and y lie in-plane. A first complication is that the precise location of the image plane (and hence the origin) within the camera is unknown, making direct depth-accuracy assessment problematic. A second complication is that small rotations or translations of the camera rig change the orientation of the coordinate frame, thereby reducing the reliability of overly simple accuracy tests.

Test setup

To address these issues, a dedicated accuracy test was constructed. The stereoscopic rig with baseline B is positioned at an axial distance $\Delta z \approx 6$ m from a rigid, planar marker target, as shown in Fig. 4.16.

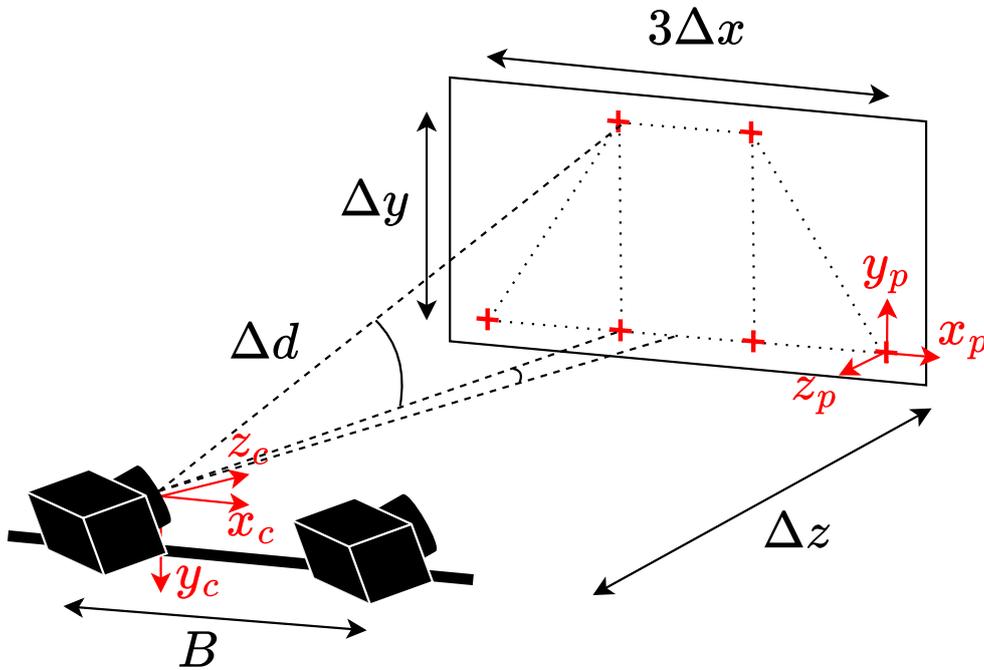


Figure 4.16: Schematic of the accuracy test. Markers (red crosses) are placed on a flat wall at a distance $\Delta z \approx 6$ m. Markers are divided by even spacings Δx and Δy . The geometry reflects the characteristic camera–marker distances in the experiments.

Markers are arranged in a trapezium on a flat wall with horizontal spacing Δx and vertical spacing

Δy . This trapezium is a simplified version of the V3 wing shape, with the maximum chord c_{\max} at the centre, decreasing towards the tips. Matching rectified frames of this trapezium structure on the flat wall are provided in Fig. 4.17. All diagonals were measured and equal, to avoid any uncertainties of a possibly uneven floor.

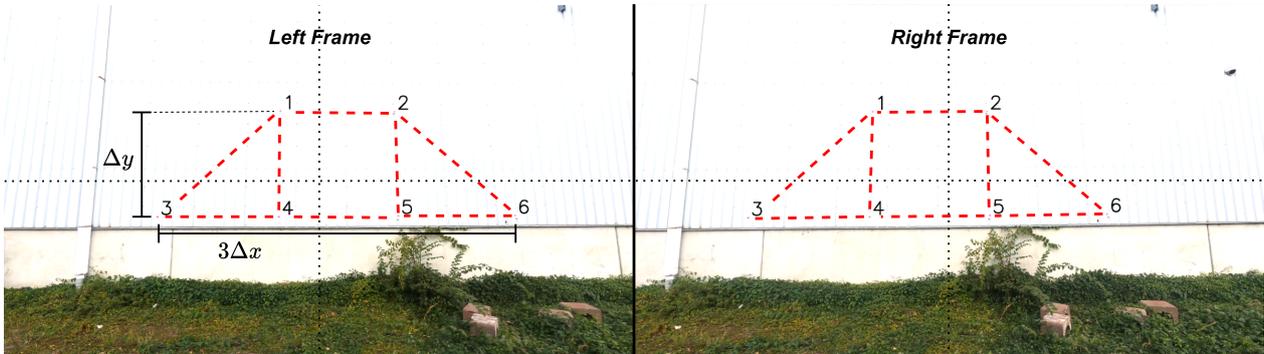


Figure 4.17: Coupled, rectified frames of the accuracy test. Identified markers are shown and numbered; the trapezium shape is illustrated for clarity.

The spacings are selected to be representative of the characteristic camera–marker distance Δd , which reaches up to about 7 m in the experiments. This distance follows from Pythagoras’ theorem applied in 3D:

$$\Delta d = \sqrt{\Delta x^2 + \Delta y^2 + \Delta z^2}.$$

The chosen geometry reflects the typical camera–wing separation during the experiments. The values are provided in Table 4.11.

Table 4.11: Metrics used in the accuracy tests described in Fig. 4.16.

Accuracy test size	
Δx	2.4 m
Δy	2.2 m
Δz	≈ 6.0 m
Δd_{\max}	≈ 7.0 m

For this rigid target, matching of markers in frames is performed grid-wise based on horizontal and vertical position. For the kite-in-flight experiments—with more markers and uneven distributions—a dedicated algorithm is used (Section 4.6). The simpler matching used here does not influence the accuracy of this rigid test.

The test yields six markers in a trapezium with reconstructed coordinates $(x_{n,c}, y_{n,c}, z_{n,c})$ in the camera frame. The markers are known to lie on a trapezium of size $3\Delta x \times \Delta y$. A plane with these dimensions is fitted to the reconstructed points via least squares. The plane is then rigidly transformed so that $z_p = 0$ everywhere on the fitted plane, allowing the points to be expressed in the plane frame as $(x_{n,p}, y_{n,p}, z_{n,p})$. Distances from each reconstructed marker to its corresponding ideal location on the fitted plane (with the real spacings) are computed.

Because the fitted plane enforces the real-world size and spacing, the residuals serve as an estimate of measurement accuracy. The error is decomposed into in-plane and out-of-plane components, $\epsilon_{p,ip}$ and $\epsilon_{p,op}$, respectively. The in-plane error quantifies the deviation within the x_p – y_p plane whereas the out-of-plane or depth error—directly related to disparity and thus to the baseline (see Fig. 4.8)—is obtained from the point-to-plane normal distance. The procedure is repeated multiple times, after which average and maximum values are reported for the Euclidean, in-plane, and out-of-plane errors.

Accuracy results for $\Delta d_{\max} \approx 7$ m

The achievable accuracy depends on the stand-off distance and the baseline; here we consider the case $\Delta d_{\max} \approx 7$ m, the expected maximum camera–marker distance. Figure 4.18 presents a representative example. The red area indicates the fitted plane (translated to coincide with the x – y plane), with ideal grid points spaced by Δx horizontally and Δy vertically.

The maximum and average errors for the in-plane, out-of-plane, and Euclidean metrics are summarised in Table 4.12.

Overall, the results meet the accuracy-related requirements (REQ-EXP-01a & REQ-DPV-01) and are consistent with expectations for this baseline, supporting the validity of the test. It is further assumed that the accuracy observed here transfers to the in-flight configuration, since the setup and viewing geometry (marker sizes, distances, and near-normal orientation to the stereo pair) are comparable. The observed maxima across repetitions indicate good repeatability and low uncertainty, underpinning the reliability of the accuracy assessment.

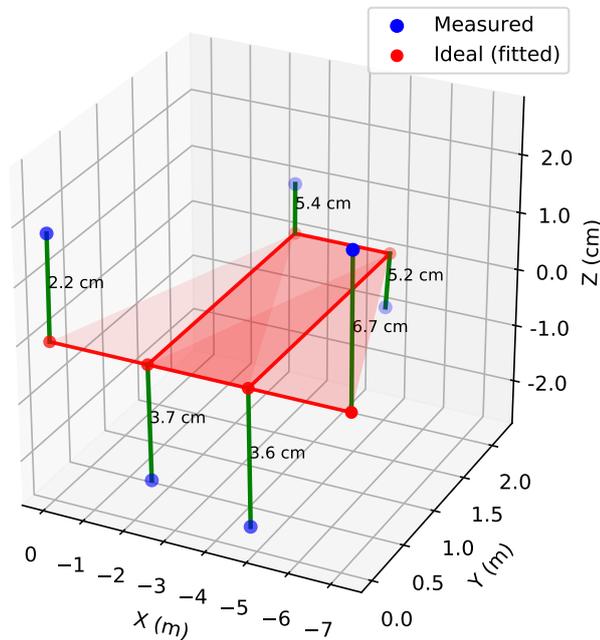


Figure 4.18: Example grid of reconstructed markers (blue) and their corresponding points on the fitted trapezium of size $3\Delta x \times \Delta y$ (red). Values shown are Euclidean errors ϵ . Note that the z axis shows centimetres.

Table 4.12: Errors for the stereoscopic setup at $\Delta d_{\max} \approx 7$ m.

Error	Value (cm)
$\epsilon_{p,ip,avg}$	3.85
$\epsilon_{p,ip,max}$	7.33
$\epsilon_{p,op,avg}$	1.34
$\epsilon_{p,op,max}$	3.89
$\epsilon_{p,avg}$	4.22
$\epsilon_{p,max}$	7.34

4.6 Photogrammetry algorithm

To obtain reliable 3D points, markers must be (i) detected in each image, (ii) correctly matched between the left and right frames, and (iii) consistently tracked over time. The pipeline below was implemented to achieve this, with epipolar constraints enforcing geometric plausibility and

lightweight tracking maintaining identities between re-initialisations. These steps are repeated for each frame independently.

Step 1: Image rectification

As detailed in Section 4.3.2, stereo rectification aligns epipolar lines horizontally, so true correspondences share (nearly) the same image row [126, Ch. 10–12]. Prior to the rectification, the undistorted image coordinates (x_u, y_u) are to be determined using the relations described in Section 4.3.1.

The calibrated intrinsics and extrinsics are subsequently used to compute rectified projection matrices, then undistort–remap each frame. The extrinsics require slight adjustments in post-processing at this step, extensively discussed in Chapter 8.

Step 2: Image adjustment

Before detection, frames are mildly enhanced to make markers stand out from the background. We apply a linear intensity transform

$$I_{\text{out}} = \text{saturate}(\alpha_{\text{image}} I_{\text{in}} + \beta_{\text{image}}),$$

using OpenCV’s `convertScaleAbs` and $\alpha_{\text{image}} = 4$ and $\beta_{\text{image}} = 20$, as used in this research:

```
1 cv2.convertScaleAbs(src, alpha=4, beta=20)
```

Here, α increases contrast (gain) and β raises brightness (bias). With these settings the red crosses are more separable under varied lighting while suppressing low-level noise [127]. Care is taken to avoid excessive clipping. The result is provided in Fig. 4.19, with the left showing the original image and the right the image with increased brightness and contrast.



Figure 4.19: Side-by-side comparison of original image (left) and image with adjusted brightness and contrast (right). With increased brightness and contrast, features on the kite become much more visible, whilst noise is removed as the background looks almost transparent.

During certain manoeuvres, such as turns, the illumination across the wing varied significantly. To compensate for this, a frame-wise adaptive brightness and contrast correction was applied. For each frame, the average brightness on both sides of the wing was computed, after which one side was classified as dark and the other as light. A brightness and contrast adjustment was then applied using a left-to-right gradient, smoothly correcting the illumination differences across the wing.

Step 3: Marker detection

Red crosses are segmented in HSV (two hue bands as red wraps around the hue axis). Contours are extracted from the processed binary image using [128]:

```
1 cv2.findContours()
```

Initially, it was hypothesized edge detection would be superior when using cross markers. Due to the limited visibility and variable lighting, edge detection failed occasionally, which is why contour detection was implemented. To eliminate false detections, each contour's area and aspect ratio is evaluated to fall within an empirically determined range, imposing a limit on the size and shape a marker can be. The centroid of each valid contour is computed using image moments:

$$(ce_x, ce_y) = \left(\frac{M_{10}}{M_{00}}, \frac{M_{01}}{M_{00}} \right),$$

where M_{ij} are the spatial moments of the contour. The resulting list of centroids corresponds to the detected marker positions.

Step 4: Setting epipolar constraints

Epipolar geometry forms the geometric foundation for establishing correspondences between points in stereo images. In an ideal stereo setup, where the cameras are perfectly calibrated and rectified, any 3D point projected onto the image planes will result in two corresponding points lying on the same epipolar line. This significantly simplifies the correspondence problem by reducing the search from a two-dimensional area to a one-dimensional line.

Mathematically, the epipolar constraint, also known as the correspondence condition [126, Ch. 9] is defined as:

$$\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0$$

where \mathbf{x} and \mathbf{x}' are the homogeneous coordinates of corresponding points in the left and right images, respectively, and \mathbf{F} is the fundamental matrix encapsulating both the intrinsic and extrinsic parameters of the stereo setup.

In practice, however, exact epipolar correspondence is rarely achieved due to residual calibration errors, noise, and imperfect rectification. Therefore, a small deviation from the perfect constraint must be tolerated. After stereo rectification, the epipolar lines are aligned horizontally, and hence the corresponding points are expected to lie on the same horizontal scanline, i.e., they should have the same or nearly the same vertical coordinate.

In the algorithm, bounded epipolar constraints are imposed by only considering matches between points in the left and right images that satisfy:

$$|\Delta y_p| < \delta_{y,\max}, \quad \delta_{x,\min} < \Delta x_p < \delta_{x,\max}$$

where $\delta_{y,\max}$ is a small vertical tolerance, and $\delta_{x,\min}$ and $\delta_{x,\max}$ are the minimum and maximum allowed disparity in the horizontal direction (depending on the distance z to the object). These bounds define a search window or bounding box around each left image point within which potential right image correspondences are considered. Δx_p is thus not symmetric, as a marker in the right frame should always be on the left side of the same marker in the left frame, after rectification. The values were tuned manually, and provided in Table 4.13. The vertical tolerance is rather high compared to similar photogrammetry studies, mainly as a result of corrections required in post-processing due to bar bending and twist (see Chapter 8). The variable horizontal disparity is necessary as the markers are located at variable distance of the camera, which alters disparity.

Table 4.13: Constraining limits used in marker matching. These values create a bounding box based on a marker in the left frame, stating that the corresponding marker in the right frame should be positioned in this box. The vertical tolerance $|\delta_{y,\max}|$ is set rather high as a result of bar bending, discussed in Section 8.2.1.

Metric	Symbol	Value (px)
Vertical tolerance	$ \delta_{y,\max} $	10
Min horizontal disparity	$\delta_{x,\min}$	200
Max horizontal disparity	$\delta_{x,\max}$	600

Step 5: Constructing a cost matrix

To formalise this further, a cost matrix \mathbf{C} is constructed of size $N \times M$, where N and M are the number of detected points in the left and right images, respectively. Each entry \mathbf{C}_{ij} is the Euclidean distance (in pixels) between point i in the left image and point j in the right image, only if it falls within the allowed disparity constraints:

$$\mathbf{C}_{ij} = \begin{cases} |\mathbf{x}_i - \mathbf{x}'_j|, & \text{if } |\Delta y| < \delta_{y,\max} \text{ and } \delta_{x,\min} < \Delta x_p < \delta_{x,\max}, \\ \infty, & \text{otherwise.} \end{cases}$$

This effectively penalises matches that violate the epipolar geometry by assigning them an unreasonably high cost, preventing them from being selected. The result is a matrix which essentially provides the correspondence between points. A visual example of the terms and principles describing the epipolar constraints and cost matrix reconstruction are provided in Fig. 4.20.

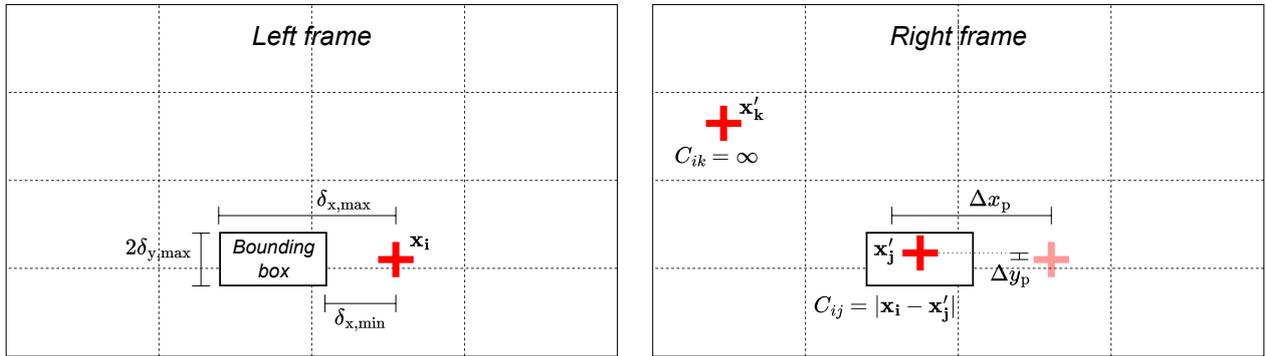


Figure 4.20: Illustration of epipolar constraints and cost matrix construction. For a detected marker in the left frame (\mathbf{x}_i), a bounding box is defined by the vertical tolerance $\delta_{y,\max}$ and horizontal disparity limits $\delta_{x,\min}$ and $\delta_{x,\max}$. Candidate matches in the right frame are only considered if they fall inside this bounding box. A marker within the box (e.g. \mathbf{x}'_j) is assigned a finite cost in the matrix equal to its Euclidean distance from \mathbf{x}_i , whereas a marker outside the box (e.g. \mathbf{x}'_k) is penalised with an infinite cost. This procedure ensures that only geometrically consistent correspondences are retained for stereo matching.

Step 6: Apply the Hungarian algorithm

The Hungarian algorithm, also known as the Kuhn-Munkres algorithm, solves the assignment problem by finding the minimum-cost matching between two sets—in this case, markers in the left and right images [129]. The goal of the Hungarian algorithm is to find a binary assignment matrix $\mathbf{A} \in \{0, 1\}^{N \times M}$ such that:

1. Each marker in the left image is matched to at most one marker in the right image, and vice versa.
2. The total cost is minimised:

$$\min_{\mathbf{A}} \sum_{i=1}^N \sum_{j=1}^M \mathbf{A}_{ij} \mathbf{C}_{ij},$$

subject to:

$$\sum_{j=1}^M \mathbf{A}_{ij} \leq 1 \quad \sum_{i=1}^N \mathbf{A}_{ij} \leq 1.$$

This combinatorial optimisation problem is efficiently solved by the Hungarian algorithm with a worst-case time complexity of $\mathcal{O}(n^3)$ for square cost matrices, where $n = \max(N, M)$. If $N \neq M$, the matrix is filled with dummy entries (with cost ∞) to make it square. The results are the matched markers, based on minimized cost.

Step 7: Filtering of matches

After the optimal matching is computed, filtering is applied to reject correspondences whose total matching cost exceeds a defined maximum threshold. This step ensures that only geometrically plausible matches are kept, improving the robustness of the system to noise, occlusions, or false detections. The result of matched markers and their epipolar lines for one time instance is provided in Fig. 4.21. Here it is clearly visible that not all markers have been detected, both as a result of (partial) occlusions as well as lighting. The frame also shows a falsely identified marker (match 38), which is removed in post-processing (see Chapter 8).



Figure 4.21: Matched markers and their epipolar lines in both the left and right frame after filtering.

Step 8: Splitting LE and strut markers

After filtering, each matched image marker still needs to be classified as either a LE or strut marker. This separation is performed directly on the image data, before triangulation, based on the design choice of a white background on the LE and a black background on the struts.

For each detected cross centroid $ce = (ce_x, ce_y)$, the local background brightness is estimated in an annular region around the marker in the value channel $V(x, y)$ of the HSV image. The annulus is defined as

$$\mathcal{A}_i = \{(x, y) \in \mathbb{Z}^2 : r_{\text{in}} \leq \|(x, y) - ce\|_2 \leq r_{\text{out}}\},$$

with inner and outer radii provided in Table 4.14

Table 4.14: Inner and outer radius used for background colour definition.

Radius	Value (px)
r_{in}	6
r_{out}	18

The median background brightness ν_i around marker i is then:

$$\nu_i = \text{median}\{V(x, y) : (x, y) \in \mathcal{A}_i\}.$$

Markers on the bright LE cloth have a high ν_i , whereas markers on the black struts have a low ν_i . Using a fixed brightness threshold ν_{thr} (tuned once per setup), each marker is first classified as

$$\text{label}(ce) = \begin{cases} \text{LE}, & \nu_i \geq \nu_{\text{thr}}, \\ \text{strut candidate}, & \nu_i < \nu_{\text{thr}}. \end{cases}$$

This results in a per-marker label in $\{\text{LE}, \text{strut}\}$ for every matched image point, which is subsequently used to build the LE polyline and the individual strut geometries after triangulation, for which the algorithm is discussed in Chapter 8.

Step 9: Triangulation

Afterwards, triangulation is applied on the remaining matches using the principle described in Fig. 4.2, where the disparity between corresponding left–right image points results in the depth. In practice, the full 3D coordinates (X, Y, Z) are reconstructed by combining this depth estimate with the known camera projection matrices: the baseline B , the focal length f , and the principal point coordinates (c_x, c_y) . These parameters map the image-plane coordinates (x, y) in each view back into real-world positions. The outcome is a set of 3D marker positions for each frame, which are then processed sequentially over time.

Step 10: Optical-flow tracking across frames or re-initialization

The identities of matched markers are preserved across time using a KLT (Kanade-Lucas–Tomasi) optical-flow tracker [130, 131]. KLT assumes brightness constancy over short intervals, i.e. the intensity of a small patch around a feature remains (approximately) the same from frame t to $t + 1$:

$$I(\mathbf{x}, t) \approx I(\mathbf{x} + \Delta\mathbf{x}, t + 1),$$

with $\Delta\mathbf{x}$ the 2D displacement. At 30 fps this “small motion” assumption is reasonable. A multi-scale formulation is used to handle moderate motions robustly. The tracking can be done in OpenCV using:

```
1 cv2.calcOpticalFlowPyrLK()
```

In the algorithm:

1. **Only the left view is tracked:** marker positions are propagated from frame t to $t+1$. Tracked markers with invalid status or large error are discarded.
2. **Re-detect and re-match on the right:** instead of tracking both views, the right image is freshly analysed each frame. Detections are matched to the tracked left points using the bounded epipolar window (vertical tolerance $|\delta_{y,max}|$, disparity range $\delta_{x,max} - \delta_{x,min}$). This avoids drift in the right stream and naturally recovers correspondences when markers reappear after short occlusions or noise. Additionally, marker numbering can be kept consistent in both frames.
3. **Preserve semantic labels:** after tracking the left points, LE and strut labels are reassigned on the current frame to keep the identity of each marker consistent, even if some are lost and later reacquired.

This design keeps the left stream temporally coherent while the per-frame stereo matching keeps geometry consistent. If the right view temporarily fails to detect a marker, the correspondence is rebuilt as soon as it reappears in the allowed search window, reducing long-term drift and avoiding identity swaps.

The entire algorithm results in a set of 3D coordinates (i.e. a point cloud) over time which together define the shape of the inflatable structure of the wing. This point cloud can be used to reconstruct the shape of the wing in post-processing (see Chapter 8), such that the shape of the wing in different flight cases can be quantitatively determined, done in Chapter 9, the results of the research.

Complementary to the stereoscopic photogrammetry setup, omnidirectional (REQ-INS-07c) UWB modules are used for ranging. This additional set of sensors will help improve the photogrammetry results through sensor fusion. This chapter will discuss the setup of this sensor system, such as the principles, sensors, calibration and range tests. This chapter first discusses the basic principle of UWB ranging in Section 5.1 and the sensors used in Section 5.2. Afterwards, it discusses both the calibration and range of the sensors in Section 5.3 and Section 5.4 respectively.

5.1 Basic principle

The UWB modules in this setup will be used for active ranging, measuring the distance between two sensors positioned on the tips of the wing, effectively determining the span b of the kite, as illustrated in Fig. 5.1. This configuration consists of a transmitter and a receiver that use ToF radio waves to calculate the distance between them.

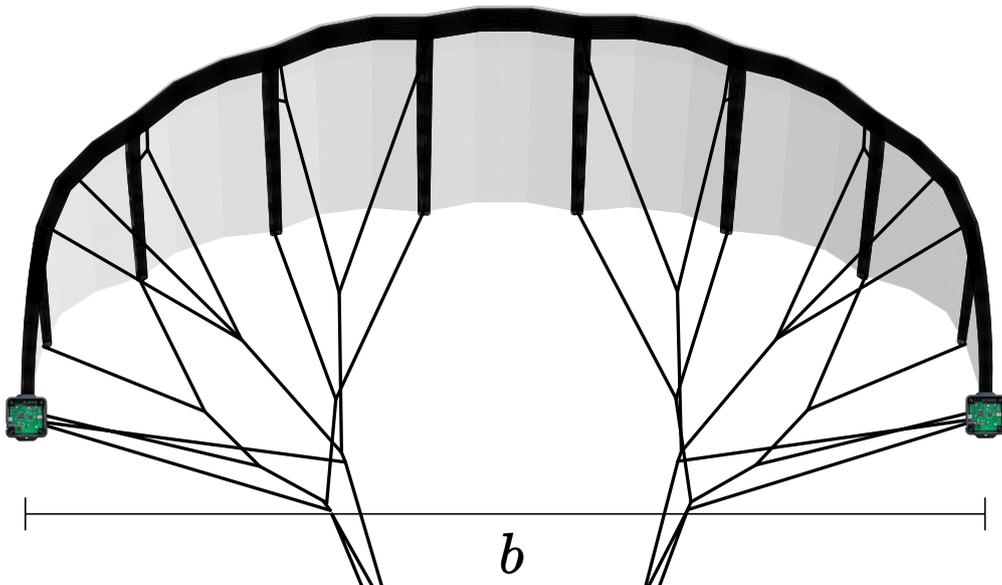


Figure 5.1: UWB sensors placed on the tips of the wing, measuring the span b .

The basic principle behind the measurement is based on the ToF method [57], where the sensors emit radio waves at the speed of light c ($3.0 \cdot 10^8 \text{ ms}^{-1}$). By measuring the time t it takes for the signal to travel from the transmitter to the receiver, the distance d between the sensors can be determined using the equation:

$$d = ct.$$

Here, d represents the distance between the sensors, and t is the time taken for the signal to travel. This distance d is thus equal to the span b of the wing.

5.2 Pozyx sensors

For the tests, a set of Pozyx¹ sensors was used. Two Pozyx anchors (V1.4) were employed for the ranging. While Pozyx tags are typically used for this purpose, the anchors were chosen for their

¹<https://www.pozyx.io/nl>

availability in a protective casing, making them better suited for the extreme conditions of the experiment. This decision also ensured robustness, as the anchors are less prone to damage in harsh environments, in accordance with the experiment’s requirements (REQ-INS-05 and REQ-INS-11).

To ensure stability, the UWB modules are securely attached to the wing tips using Velcro, with an additional safety line in place to prevent detachment in case of sudden failure. The ranging instructions for these anchors are transmitted via a Pozyx shield (v1.5), which eliminates the need for a direct cabled connection. Instead, the anchors receive instructions via the shield and operate independently. The sensors are shown in Fig. 5.2.

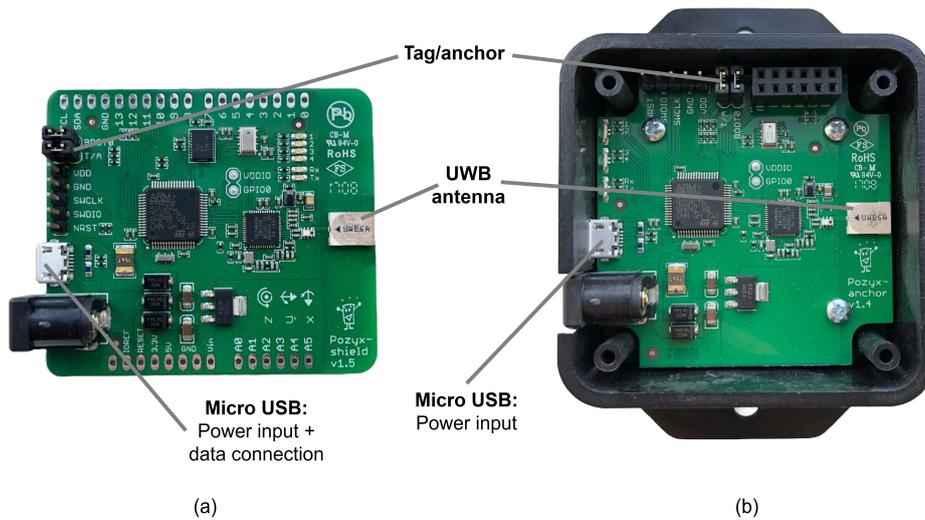


Figure 5.2: Pozyx sensors used for the experiment: (a) Pozyx shield v1.5, used for sending commands and subsequently receiving ranging data. (b) Pozyx anchor v1.4, used for ranging between one another. The relevant components of both sensors are highlighted.

The Pozyx anchors are powered by external powerbanks (5 V, 1 A), which provide sufficient energy for the system. The anchors perform the ranging and send the data back to the Pozyx shield, which is connected to a computer that collects and processes the ranging data.

Table 5.1: Relevant parameters of the Pozyx sensors. Only specifications relevant to this study are listed. Values on accuracy and range seem to differ greatly, and are therefore analysed independently for the specific set of sensors and setup.

Parameter	Value
Sampling frequency	15 Hz
Shield mass	16 g
Anchor mass	65 g

The relevant parameters of the Pozyx sensors used in this experiment are summarized in Table 5.1, including the sampling frequency and the masses of the shield and anchors. The sampling frequency f_{UWB} at 15 Hz is sufficient to capture detailed differences in the span b (REQ-EXP-04). The accuracy and range of the system were found to vary greatly, and these were analysed independently based on the specific sensor setup for the experiment.

The sensors can be controlled using Python, via the available `pypozyx` library [132], specifically designed for sensors of this brand:

```
1 import pypozyx
```

In this experiment, the command sent to the anchors is as follows:

```
pozyx.doRanging(receiving_anchor_id, device_range, transmitting_anchor_id)
```

This command allows one anchor to act as the transmitter and the other as the receiver. The distance between the two anchors, denoted as b , is then transmitted back to the shield for further processing.

5.3 Calibration

To ensure the highest accuracy, the sensors were calibrated using a linear calibration equation:

$$d_{\text{calibrated}} = a_{\text{UWB}} \cdot d_{\text{raw}} + b_{\text{UWB}}.$$

In this process, the uncalibrated data was obtained by placing the sensors at various known distances x from each other, measured using a laser distance meter with an uncertainty of 1 mm. Since this uncertainty is negligible, it was ignored, and the actual distance between the sensors was assumed to be the measured value. Distances ranging from 5 to 15 meters, in steps of 2 meters, were used, spanning more than the expected operational range of the sensors. For each distance, three 5-second samples were taken, and the average raw distance d_{raw} was calculated.

Calibration results

The calibration principle followed a linear approach, as it is commonly used in scenarios where the error is mostly due to a constant offset—often caused by factors such as delays in receptor time or signal processing discrepancies (see Chapter 2). This method minimizes the sum of squared errors (SSE) between the predicted values and the actual measured values using linear least squares regression. The resulting calibration constants a and b are listed in Table 5.2.

Constant	Value
a_{UWB}	0.9992
b_{UWB}	0.0793

Table 5.2: Calibration values for the UWB sensors

An error analysis of the calibration test is shown in Fig. 5.3. The results indicate a consistent offset in the raw measurements, with the sensors systematically underestimating the distance between the two sensors. Before calibration, the average absolute error $\epsilon_{\text{UWB,raw}}$ was 7.17 cm.

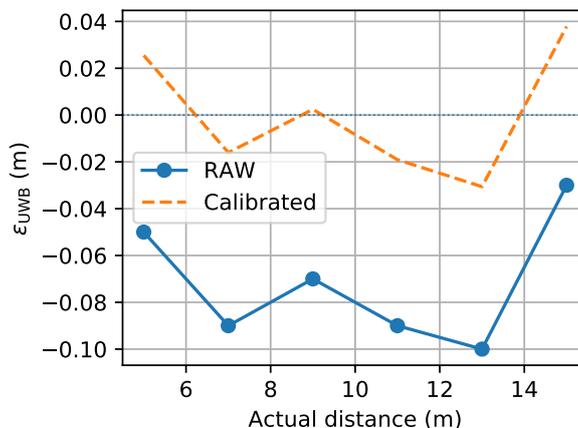


Figure 5.3: The error (ϵ_{UWB}) of the UWB sensors before and after calibration, for distances ranging from 5 to 15 meters.

After applying the calibration, the corresponding average error $\epsilon_{\text{UWB,calibrated}}$ was reduced to 2.19 cm. Studies have shown that relative movements between the sensors do not significantly affect the

accuracy of the measurements. Therefore, it is assumed that the calibrated average accuracy will be maintained during flight, though the maximum errors will differ, also as a result of sensor noise. It is furthermore assumed the environmental conditions will not have a large effect on ranging accuracy, as derived from the literature review [59, 60]. This should be valid as the sensors are close to one another, meaning interruption is minimal. Furthermore, weather conditions are expected to remain similar throughout the short duration of the experiments.

5.4 Range test

To assess whether the Pozyx shield could be positioned on the ground during experiments or should instead be flown with the kite - complicating the setup - the effective communication range of the UWB system was tested. The evaluation was performed in open-air conditions to replicate the expected experimental environment. Multiple UWB bandwidths were tested, but no significant difference in performance was observed. To do this, the first radius r_1 Fresnel zone of the setup was evaluated [133], detailed in Appendix C.

The measured range is shown in Fig. 5.4. In open-air conditions, the signal dropped after approximately 25 m, well below the required 400 m as well as values in literature which show performance at over hundred metres in certain experiments (see Chapter 2). This result demonstrates that ground-based placement of the shield is not feasible, as communication with airborne anchors would be lost. Consequently, the UWB shield, including the computer, must be flown with the kite to remain within range of the anchors.

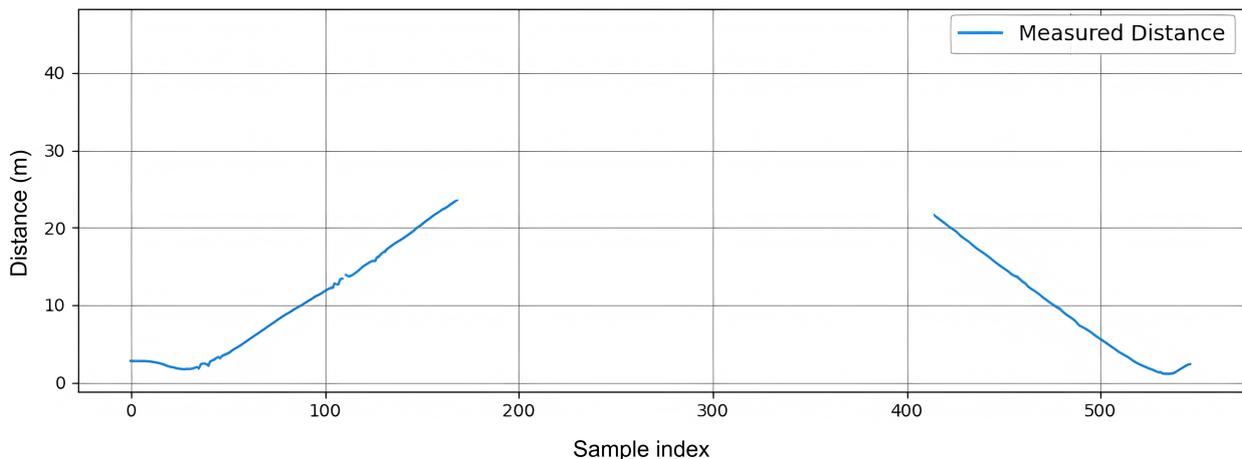


Figure 5.4: Range test between a Pozyx anchor (v1.4) and Pozyx shield (v1.5). The communication link was lost at ≈ 25 m, insufficient for ground-based operation.

Improving data continuity

Although the two anchors mounted on the kite tips consistently remained within each other’s communication range, the ground-based shield exceeded the effective transmission limit during flight. To ensure uninterrupted data acquisition, the shield was therefore integrated into the airborne system and mounted in proximity to the kite.

A Raspberry Pi 3A+ was used to execute the Python ranging code and to store the data locally. The unit, together with the shield, are powered by a compact power supply (5 V/2.5 A) and could be remotely accessed via SSH for initialization and monitoring. This configuration ensured that the shield remained within ≈ 10 m of both anchors at all times (REQ-INS-07a), thereby maintaining reliable communication and enabling continuous span measurements. The Raspberry Pi model used in the experiments is shown in Fig. 5.5, indicating the relevant components used for this research.

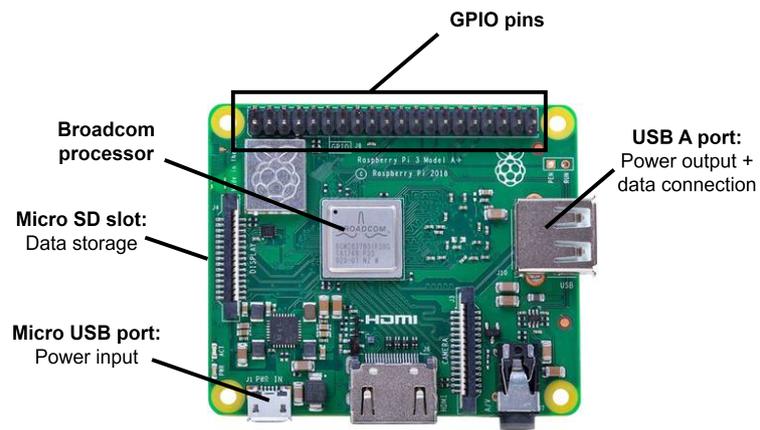


Figure 5.5: Raspberry Pi 3A+ used to run the UWB ranging code and store data during flight. Only relevant subsystems have been highlighted.

This stream and sensor layout together comprise the UWB subsystem, used as complementary data to the photogrammetry results.

Wind speed & angular measurements 6

Wind speed and angular measurements are essential for ensuring the reproducibility of tests and for evaluating the effects of wind speed on the kite's behaviour (REQ-INS-09 & REQ-DPV-03c). This chapter first discusses the principle, calibration and sensitivity of the Pitot tube used for wind speed measurements in Section 6.1, whereafter it discusses the principle and calibration of the angular vanes Section 6.2.

6.1 Apparent wind speed measurements

This section discusses the apparent wind speed measurements using a pitot tube, capturing the principle, calibration and directional sensitivity. A Pitot tube system was chosen for this purpose, using technology from Kitepower B.V. [11]. The exact model of the Pitot tube is unknown, but its basic principle and the associated data acquisition system are standard technologies for measuring wind speed in aerodynamics research. The pitot tube was found to have a sampling frequency close to $f_{ws} \approx 11$ Hz (REQ-EXP-04). Additionally, the casing around the Pitot tube houses angular measurement devices, capable of measuring angles along two rotational axis.

Basic principle

The Pitot tube operates by comparing the static and dynamic pressures. The total pressure p_{tot} is the sum of the static pressure p_s and the dynamic pressure $p_{dynamic}$, which is related to the wind speed. The dynamic pressure can be expressed using Bernoulli's principle:

$$p_{dynamic} = \frac{1}{2}\rho V_{\infty}^2 = p_{tot} - p_s,$$

where ρ is the air density and V_{∞} is the free stream velocity. By measuring the total pressure and subtracting the static pressure, the wind speed can be determined as follows:

$$V_{\infty} = \sqrt{\frac{2(p_{tot} - p_s)}{\rho}}.$$

A diagram of the Pitot tube principle is shown in Fig. 6.1, where the free stream velocity V_{∞} flows into the tube, causing a measurable total pressure at the front, while the static pressure is measured perpendicular to the flow.

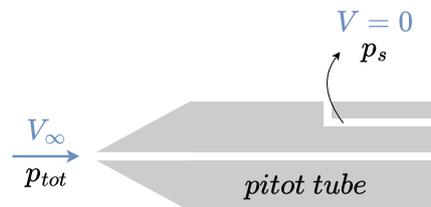


Figure 6.1: Pitot tube principle. The free stream velocity V_{∞} flows in the front, leading to a total pressure p_{tot} measured there. Perpendicular to the flow, the flow velocity is equal to $V = 0$ ms^{-1} , with static pressure p_s .

6.1.1 Experimental setup

The accuracy and sensitivity of the Pitot tube was evaluated in the M-tunnel at TU Delft [134]. The wind tunnel was operated in an open configuration whilst turbulence was low due to its large

contraction ratio. The test section of the wind tunnel has a square cross-section of $0.4\text{ m} \times 0.4\text{ m}$. Tests were conducted both inside and outside the wind tunnel, in an attempt to simulate different inflow angles and boundary layer effects. A schematic of both configurations is provided in Fig. 6.2, where V_a is the magnitude of the "apparent wind speed", which is measured by the pitot.

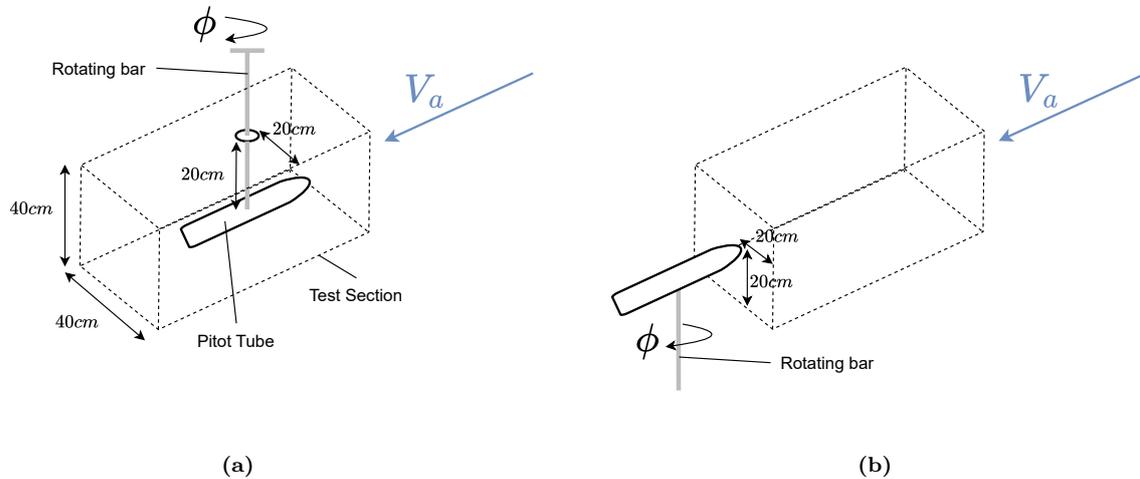


Figure 6.2: Wind Tunnel test setup, with measurements conducted (a) inside and (b) outside the wind tunnel. Here, V_a is the magnitude of the measured apparent wind speed.

The sensor casing was mounted on a rotating bar to adjust the inflow angle precisely, as depicted in Fig. 6.3. During initial tests, angular vanes on the casing did not respond at low wind speeds, and misalignment due to turbulence was observed. This led to the adoption of a protractor-based method for measuring the inflow angle more accurately, as shown in Fig. 6.4.



Figure 6.3: The experimental setup inside the wind tunnel. The casing (in white) has two vanes for angular measurements and the Pitot tube on the left.

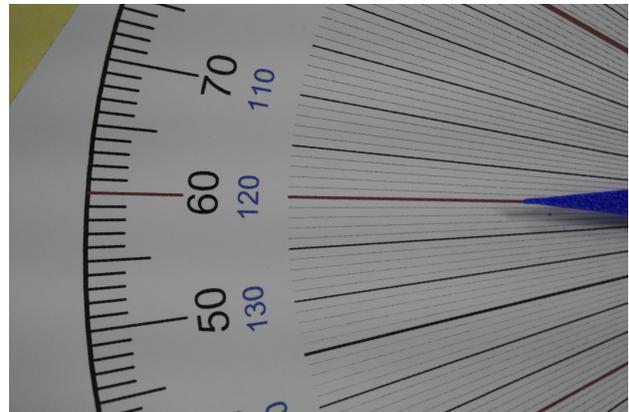


Figure 6.4: Protractor and 3D-printed indicator used to accurately determine the inflow angle ϕ , mounted on the rotating bar.

Wind speed measurements were taken across a range of speeds and inflow angles. Each test was assigned a unique identifier to facilitate data processing. Measurements were averaged over a period of 5 seconds, corresponding to roughly 55 data points, given the 11 Hz sampling frequency of the Pitot tube.

6.1.2 Calibration and directional sensitivity

Calibration

To ensure the accuracy of the wind speed measurements, the Pitot tube was first calibrated at an inflow angle of $\phi = 0^\circ$ for different flow speeds, where the tube is directly aligned with the

airflow. The measured velocity should match the actual flow speed at this alignment. The calibration relationship is given by the linear equation:

$$V_{\text{calibrated}} = a_{\text{ws}} \cdot V_{\text{raw}} + b_{\text{ws}}$$

The calibration constants a and b were determined through linear regression at $\phi = 0^\circ$, and are provided in Table 6.1. This calibration model was then applied to all subsequent measurements.

Constant	Value
a_{ws}	1.064
b_{ws}	-0.415

Table 6.1: Calibration constants used for the Pitot tube.

Directional Sensitivity

The Pitot tube's directional sensitivity was evaluated by measuring the wind speed for different inflow angles. As seen in the results shown in Fig. 6.5, accuracy declines at inflow angles beyond approximately 25° to 30° , particularly at higher wind speeds. However, small variations between adjacent angles suggest that minor misalignments between the kite and the sensor do not significantly impact the accuracy of the measurements.

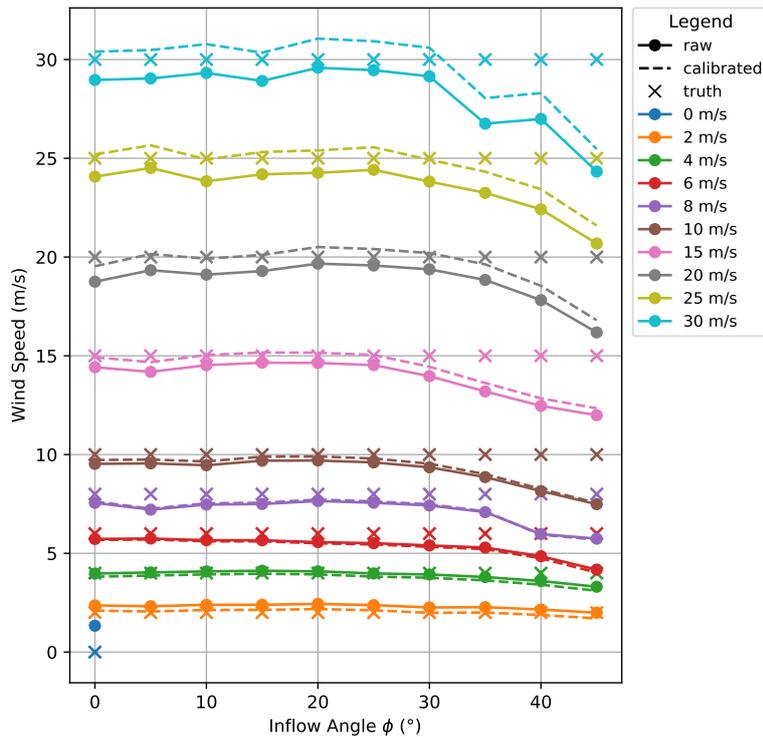


Figure 6.5: Measured wind speed for different inflow angles ϕ before and after calibration.

The average absolute and relative errors at various inflow angles are provided in Table 6.2. The average error across all wind speeds was calculated for each inflow angle, following the equation:

$$\epsilon_{\text{ws,avg}} = \frac{1}{N} \sum_{i=1}^N |\bar{V}_{\text{measured},i} - V_{\text{actual},i}|,$$

where $\bar{V}_{\text{measured},i}$ is the average measured wind speed of the sample for the i -th wind speed and $V_{\text{actual},i}$ is the actual wind speed for the same measurement.

The results indicate that for inflow angles less than 35° , the errors remain below 3%, fulfilling the accuracy requirement (REQ-EXP-01). However, for angles beyond 35° , the errors increase, which necessitates careful mounting of the Pitot tube to maintain measurement accuracy within the required range.

Table 6.2: Average absolute and relative error for the Pitot tube.

Inflow angle ϕ ($^\circ$)	$\epsilon_{ws,avg}$ (ms^{-1})	$\epsilon_{ws,avg}$ (%)
0	0.340	2.831
5	0.343	2.569
10	0.263	1.972
15	0.228	1.708
20	0.359	2.694
25	0.372	2.788
30	0.369	2.768
35	0.823	6.174
40	1.404	10.534
45	2.412	18.091

These results suggest that to maintain a maximum uncertainty of 1 ms^{-1} , the inflow angle should not exceed approximately 35° . No significant error trends were observed at lower angles, as errors remained consistently below 3% for angles up to 30° , ensuring sufficient accuracy within the operational range. It is assumed that the angle between the Pitot tube and the wing remains below $\pm 35^\circ$, which is guaranteed by the high tension in the front bridle lines.

6.2 Angular measurements

This section discusses the angular measurement sensors, which were redesigned to better capture these angles throughout the process.

Basic Principle

Two free-weather vanes are mounted on the casing: one sensitive to the vertical component of the flow (proxy for α) and one to the lateral component (proxy for β_s). Each vane aligns with the local streamlines (ϕ) and its angle is read by an encoder. In motionless air the vane is statically balanced to zero torque; in flow, the aerodynamic moment drives it to an equilibrium where the vane chord is (nearly) parallel to the local velocity vector. The measurement is therefore a direct, quasi-steady estimate of inflow angle at the vane location. The working principle is depicted in Fig. 6.6. In the process, the CAD design of these 3D-printed angular vanes were improved to better capture the inflow angle ϕ , further elaborated in Appendix D.

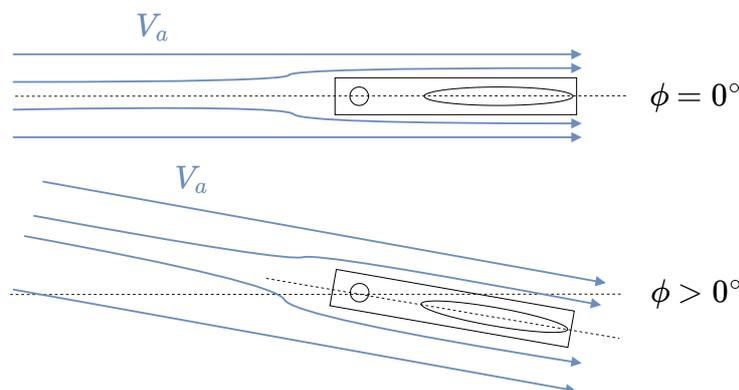


Figure 6.6: Working principle of the angular vanes attached to the pitot casing.

6.2.1 Calibration

The test setup mirrors the pitot calibration: the assembly is rotated to commanded angles indicated by the protractor, while the vane encoder output is averaged over the first 5 s of each run. A simple global linear model is used,

$$\phi_{\text{calibrated}} = a_{\text{av}} \phi_{\text{raw}} + b_{\text{av}},$$

with coefficients obtained by least-squares over the subset with reliable vane response ($V_a \geq 10 \text{ ms}^{-1}$). The resulting constants are listed in Table 6.3 and applied to all points. As shown in Fig. 6.7, calibration removes most of the large-angle overshoot while preserving near-zero behaviour.

Constant	Value
a_{av}	0.9195
b_{av}	0.4202

Table 6.3: Calibration constants used for the angular vanes.

The result of calibrated vs uncalibrated results is provided in Fig. 6.7. One can see that the uncalibrated results still show an overshoot of the estimated angle due to the vane being in the wake of the casing, particularly at higher inflow angles and wind speeds. This overshoot is significantly reduced through the calibration.

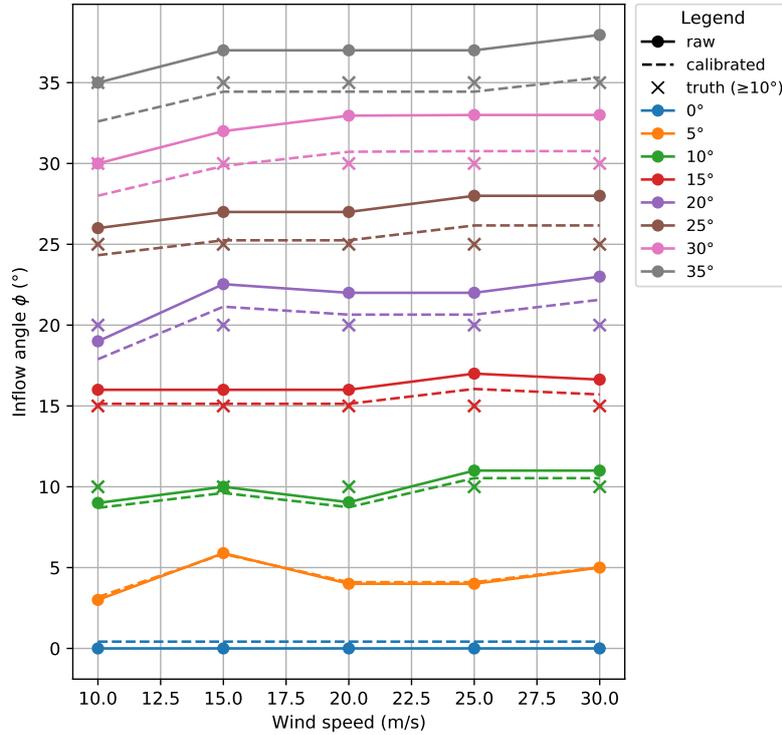


Figure 6.7: Angular vane measurements before and after calibration. The calibration clearly improves the results of the angular measurements, which overshoots for larger angles of attack, particularly at higher wind speeds.

With a maximum error after calibration of $|\epsilon_{\text{av,max}}| = 2.3^\circ$ over the evaluated range, sufficient for the purpose of this research, where angles don't need to be exact, but rather the global change throughout flight.

Experimental design 7

This chapter details the experimental design used to acquire in-flight measurements. The experiment was done at Kitepower B.V.'s test site in Bangor-Erris, Ireland. There, a fully integrated system is in use for research and design. The chapter first describes the mechanical integration of all sensors on the kite in Section 7.1, including the rigid support (" Γ ") bar, hereinafter referred to as the Γ -bar. It then introduces the KCU together with the signals used to quantify depower and steering in Section 7.2. Afterwards, Section 7.3 provides an assessment of intrusiveness (aerodynamic blockage and added mass), followed the time-base synchronisation between the different measurement and sensory systems in Section 7.4. Finally, the chapter will end with an overview of measurement architecture (Section 7.5).

7.1 Experimental setup

The measurement payload is carried by a rigid Γ -bar suspended in the front-bridles beneath the wing, following the general placement philosophy used in prior work (e.g. [49]). Instead of a " Π " shaped bar with two vertical stabilizers, one stabilizer was deemed superior to reduce twist on the bar, lowering the change of misalignment of the cameras compared to the stereoscopic calibration. The assembly is shown schematically in Fig. 7.1. The bar is mounted in the front lines as the stability is directly related to the tension in the bridles, which is much higher in the front lines than the rear lines. The two synchronised cameras are mounted at the ends of the bar with a fixed stereo baseline B . Their optical axes are directed toward the lower surface of the wing, which carries all markers.

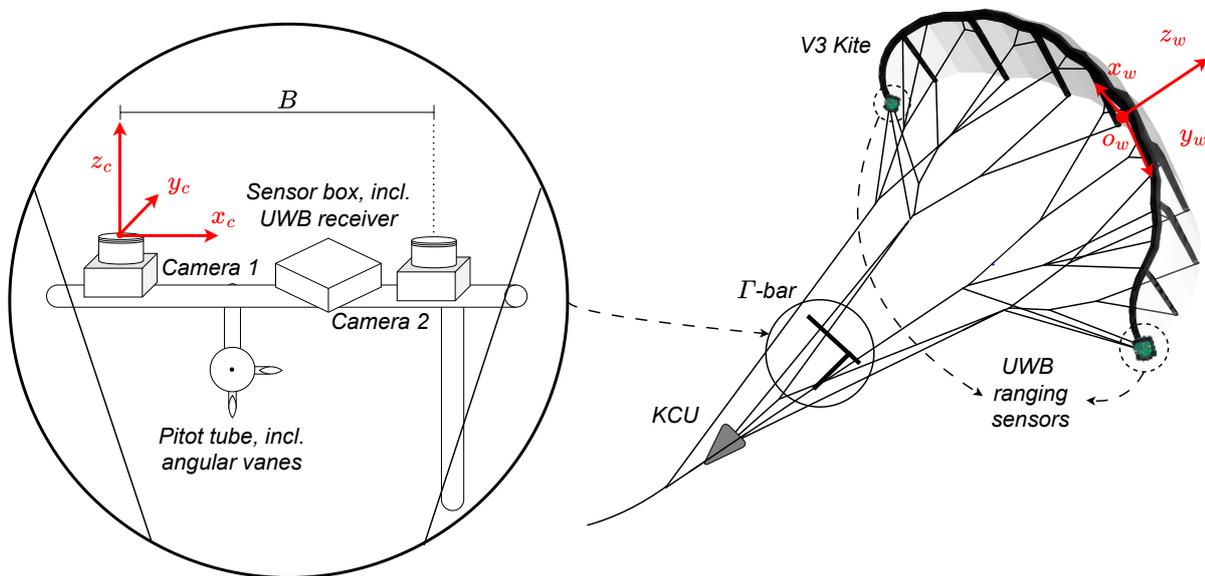


Figure 7.1: Schematic of the experimental setup. A rigid Γ -bar, mounted in the bridles beneath the wing, carries two cameras separated by a stereo baseline B , a Pitot tube, and a sensor box (including the UWB receiver). UWB ranging sensors are installed at both wingtips to infer the instantaneous span b . The cameras' fields of view are directed toward the lower surface of the wing.

The Pitot tube and vanes are attached to the bridle structure near the bar, aligned with the local bridle axis to limit angular misalignment during flight. A compact sensor box mounted on the

bar houses power distribution and on-board electronics (including the UWB receiver). The UWB ranging anchors are mounted at the wingtips so that their mutual distance directly provides the instantaneous span b . Each anchor is fixed to the tip fabric using Velcro and duct tape, and secured with a short safety leash to the tip webbing to prevent loss in the unlikely event of adhesive failure. Power is supplied by compact powerbanks tucked beneath the wing’s leading-edge fabric adjacent to the inflatable bladder, minimising aerodynamic disturbance and keeping the mass close to existing structure. This configuration yields stable camera viewpoints, ensures that the fields of view cover the marker pattern throughout manoeuvres, and minimizes cable runs.

The supporting Γ -bar structure consists of carbon- and glass-fibre tubes ($\varnothing 2$ cm). This geometry introduces a vertical separation between the bridle clamps, which reduces sensitivity to vibrations and external disturbances (REQ-EXP-07, REQ-INS-02). Under flight conditions, the bridles experience high tension, which stiffens the setup and restricts relative movement. As a result, the entire wing remains within the camera field of view even at high wind speeds, fulfilling the stability requirement (REQ-EXP-04). The overall configuration in flight is shown in Fig. 7.2a, and the Γ -bar is detailed in Fig. 7.2b.

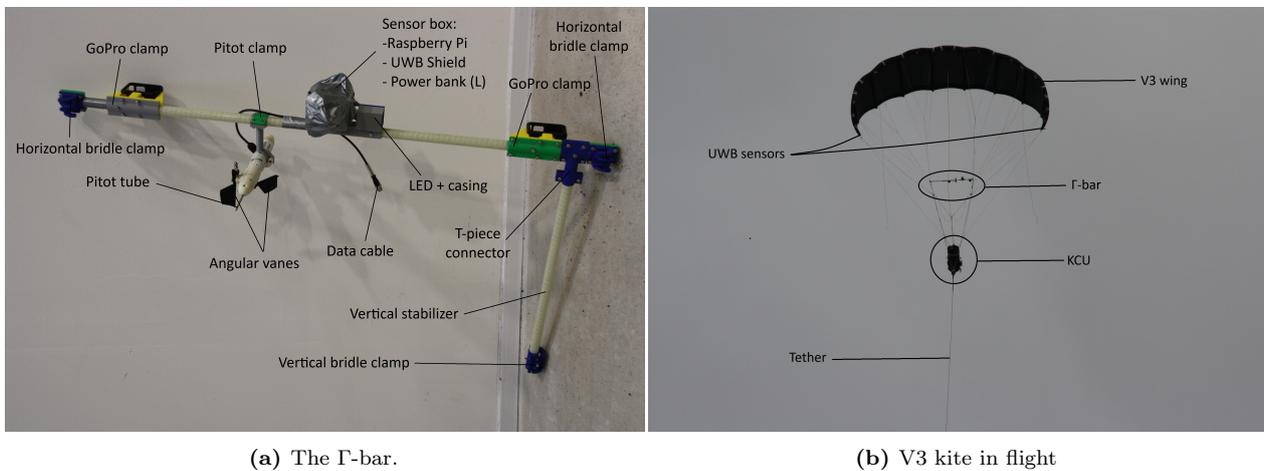


Figure 7.2: Experimental setup: (a) Γ -bar detail; (b) overall in-flight configuration.

The Γ -bar and sensor subsystems are assembled using several custom-designed 3D-printed components. A detailed description and graphical overview of these parts is provided in Appendix E. Briefly, the design objectives were to achieve stable and fixed camera viewpoints, ensure zero relative motion between the cameras, and provide reliable clamping of the bar to the front bridles. The printed components include a pitot clamp for securing the pitot tube and angular vanes, as well as a sensor box that houses part of the UWB subsystem. Together, these elements provide stable camera geometry, continuous visibility of the marker pattern, and straightforward attach-and-detachment.

Measured angles

The angular vanes mounted on the Pitot casing sense the flow direction relative to the front-line bridle, not the wing’s chord. The measured “angle of attack” is therefore the front-line inclination, α_{fl} , whereas the aerodynamic angle of attack at the wing is α_w , as depicted in Fig. 7.3. Because the front lines are long and the wing chord is short, the mapping of α_{fl} and α_w is nearly constant during flight, though absolute values differ. Depower primarily rotates the wing about its bridle attachment and can change α_w substantially without a commensurate change of α_{fl} .

For the measured sideslip, β_s , we assume it approximates the wing sideslip because, under tension, the front lines rotate with the wing and transmit lateral orientation with little elastic lag. This matches observations in the wind tunnel and in flight. These angles are additionally used to assure the Pitot remains within the specified angular range.

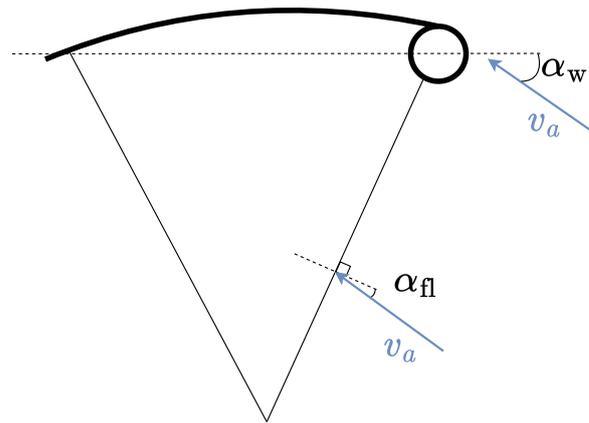


Figure 7.3: Measured α_{fl} (front-line inclination) versus aerodynamic angle at the wing α_w .

7.2 KCU

The KCU is a cable robot that actuates the wing using battery-driven motors. These motors execute preprogrammed flight patterns and enable energy generation. Two independent actuation modes are used:

- (i) **depower** — a symmetric change of the TE bridle lengths via the depower tape, which primarily changes the effective angle of attack; and
- (ii) **steering** — an asymmetric change of left/right TE bridle lengths via the steering tape, which induces yaw/roll and turns the kite.

A schematic overview of the KCU and tapes is provided in Fig. 7.4, where the dotted lines indicate bridle line connections. The lengths of bridles and all lines affected during flight are provided in

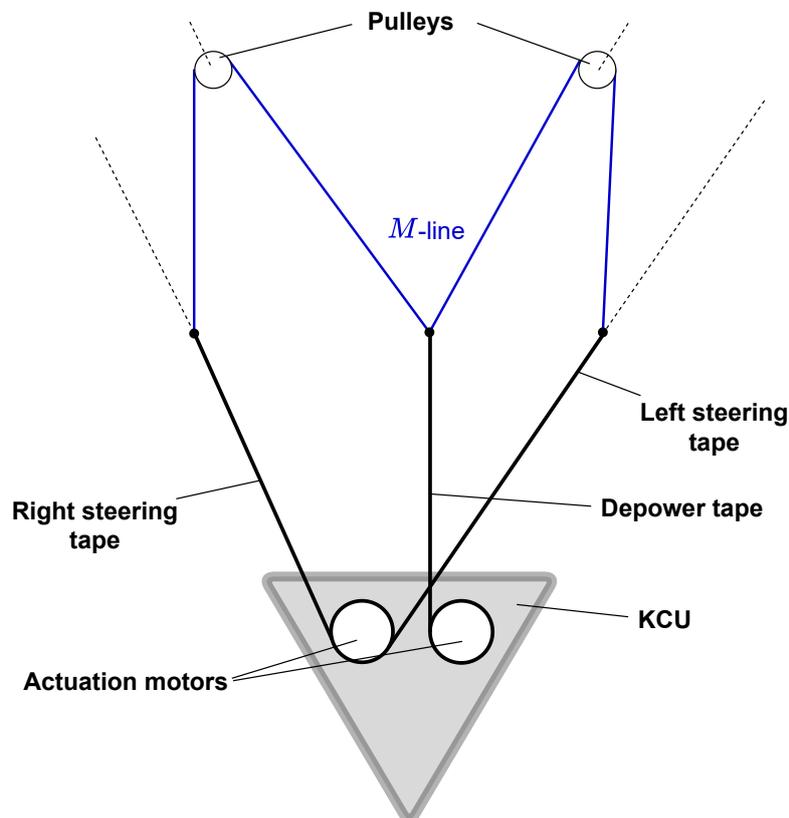


Figure 7.4: Front view of KCU actuation; depower and steering tapes through motors.

Appendix F, including their connection points on the wing.

7.2.1 Definition of inputs

KCU commands are given as dimensionless inputs: depower $u_{dp} \in [0, 1]$ and steering $u_s \in [-1, 1]$.

Power/Depower

Depower is commanded by changing the depower-tape length L_{dp} (mm). A larger L_{dp} corresponds to a more depowered wing. Based on the calibration used here, L_{dp} varies linearly with u_{dp} :

$$L_{dp}(u_{dp}) = 200 + 5000 u_{dp}, \quad u_{dp} \in [0, 1],$$

so that $L_{dp}(0) = 200$ mm, $L_{dp}(1) = 5200$ mm, and a step of $\Delta u_{dp} = 0.01$ changes tape length by $\Delta L_{dp} \approx 50$ mm. Note that u_{dp} here corresponds to the depower convention as used by Kitepower B.V. [11].

Table 7.1: Depower settings.

u_{dp} (-)	L_{dp} (mm)
0.00	200
1.00	5200
+0.01 step	+50 per step

Turning manoeuvre

Steering is commanded by a differential change between the left/right steering tapes. We define $u_s < 0$ as a left turn and $u_s > 0$ as a right turn. Let $L_0 = 1600$ mm denote the neutral tape length on each side, and let k_s (mm) be the steering gain (calibrated). A simple linear mapping is

$$\begin{aligned} L_{\text{left}}(u_s) &= L_0 - \frac{k_s}{2} u_s, \\ L_{\text{right}}(u_s) &= L_0 + \frac{k_s}{2} u_s, \quad u_s \in [-1, 1], \end{aligned}$$

so that a negative u_s shortens the right tape and lengthens the left tape (right turn), and vice versa for $u_s > 0$.

The magnitude k_s sets the maximum differential (k_s at $|u_s| = 1$) and is determined from KCU calibration and mechanical limits, the settings are depicted in Table 7.2. Note that the total length depends on the input due to differential steering, resulting in more/less tape spooled around the actuation motor.

Table 7.2: Steering settings.

u_s (-)	L_{left} (mm)	L_{right} (mm)
0	1600	1600
+1	200	2800
-1	2800	200

In practice, both u_p and u_s are rate-limited and saturated by KCU firmware and mechanical end-stops; the linear relations above describe the nominal mapping used for analysis and logging. In this research, the definition of sides of the kite in a turn will be the pulled and released side. Where for example in a left turn, the left tip is pulled, whereas the right tip is released.

7.3 Intrusiveness

This section evaluates the intrusiveness of the entire experimental setup to the kite and its behaviour through both flow aerodynamic alternation and mass.

7.3.1 Flow obstruction

The experimental setup was designed to minimise aerodynamic disturbance. The rigid bar is positioned well below the wing, ensuring that the canopy’s flow field remains essentially unaffected. The structure is deliberately placed neither upstream of the wing—where it would interfere with the incoming airflow—nor in the wake region, where it could alter vortex shedding. The bar lies far beneath the wing and its flow field, and its effect is assumed negligible.

On the wing, surface markers consist of thin adhesive tape ($th_{\text{tape}} \approx 230 \mu\text{m}$) applied directly to the canopy. While even small surface imperfections can locally alter boundary-layer behaviour (e.g. transition or separation), the magnitude of these additions is much smaller than existing surface irregularities introduced during kite manufacturing. On top of this, the boundary layer is most sensitive when it remains laminar and attached, which is not expected on the pressure side. There, the circular leading-edge geometry promotes early separation and a predominantly turbulent flow field [24].

The only larger external hardware are the UWB anchors ($th_{\text{anchor}} \approx 15 \text{ mm}$), which are mounted near the trailing edge at the wing tips. While these anchors are comparatively bulky, they remain small relative to the overall wing dimensions and are located exactly where tip vortices dominate (see Chapter 2). Experimental studies of winglet and small wingtip devices show that small-span alterations have only limited effect on the global vortex structure—particularly when placed in a region already dominated by vortex dynamics [135]. Therefore, although minor local flow perturbations may occur, the large-scale vortex behavior and overall aerodynamic performance of the wing are expected to remain intact. The UWB anchors are powered by compact powerbanks concealed beneath the leading-edge inflatable structure, ensuring the canopy geometry is undisturbed and the added mass is co-located with existing structural elements. Consequently, the aerodynamic impact of the experimental setup is considered negligible (REQ-INS-04a).

7.3.2 Mass

The total mass contribution of the experimental setup is summarised in Table 7.3. All airborne components are included, while ground-based hardware is excluded as it does not influence flight dynamics. Lightweight elements such as surface markers ($\approx 3 \text{ g/m}$ dacron tape, with only a few metres used) and small LED components are ignored due to their negligible contribution. The mass of the experimental setup is 6.93% of the kite and KCU mass.

Table 7.3 shows that the complete system has a mass of less than 2.5 kg, which is within the requirement of remaining below the mass of the airborne turbine used to provide power to the KCU (REQ-INS-04b). This makes the system suitable for application on kites of comparable or larger size. In future upscaling, only the length of the composite bar would change, with minimal impact on total mass. For smaller kites, the current setup may be too heavy; in such cases, a reduced instrumentation suite (e.g., omitting UWB sensors) could be considered. This ensures the system is easily adaptable for a range of kites (REQ-EXP-05).

7.4 Time-synchronisation

To ensure that all sensing modalities produce temporally consistent data, their internal clocks must be synchronised (REQ-DPV-08). Without synchronisation, steering and depower inputs could not be coupled to measured wing deformations, the UWB and photogrammetry datasets could not be cross-validated, and wind and attitude measurements could not be interpreted reliably. Synchronisation therefore enables each geometric deformation to be linked unambiguously to the corresponding flight condition.

Table 7.3: Mass contribution of each item of the experimental setup. Only airborne items are listed; total masses below 15 g are neglected.

Item	Mass per unit (g)	Amount	Mass (g)
Photogrammetry setup (Chapter 4)			
GoPro Hero 9 Black	158	2	316
GoPro case	60	2	120
GoPro clamp left	40	1	40
GoPro clamp right	45	1	45
UWB setup (Chapter 5)			
Pozyx Anchor V1.4	65	2	130
Pozyx Shield V1.5	16	1	16
Raspberry Pi 3A+	30	1	30
USB-A to Micro-USB cable	22	4	88
Power bank (L)	131	1	131
Power bank (S)	64	2	128
Sensor box	256	1	256
Pitot setup (Chapter 6)			
Pitot+vanes, incl. cable	172	1	172
Pitot clamp	47	1	47
Γ-bar assembly			
Horizontal GF bar	455	1	455
Vertical GF bar	105	1	105
Bar connector	30	1	30
T-Piece connector	35	1	35
Horizontal bridle clamp	31	2	62
Vertical bridle clamp	28	1	28
Bolts (M4x16)	2	56	112
Nuts (M4)	0.5	56	28
Total			2374

7.4.1 Camera synchronisation

Accurate stereo triangulation requires that the left- and right-camera frames correspond to the same physical instant. The GoPro Hero9 Black does not provide hardware synchronisation, and its internal clocks are too inaccurate for stereo work: offsets of centiseconds to seconds are common, and drift accumulates over longer recordings. Such discrepancies easily exceed the frame period ($\approx 1/30$ s) and would degrade reconstruction quality.

For this reason, synchronisation is performed entirely in post-processing by aligning the audio streams recorded by both cameras. The GoPro audio is sampled at 48 kHz, providing more than sufficient temporal resolution for sub-millisecond alignment. It is assumed that both cameras have the same internal audio delay and that the sound source used for synchronisation—hand claps—was positioned approximately equidistant to both cameras.

Two timing components must be accounted for, depicted in Fig. 7.5:

1. lag — the offset between the start times of both recordings, and
2. drift — gradual desynchronisation of the clocks over time.

The frame period is

$$T_{\text{frame}} = \frac{1}{fps},$$

and the initial lag is

$$\text{lag} = t_{\text{start},2} - t_{\text{start},1}.$$

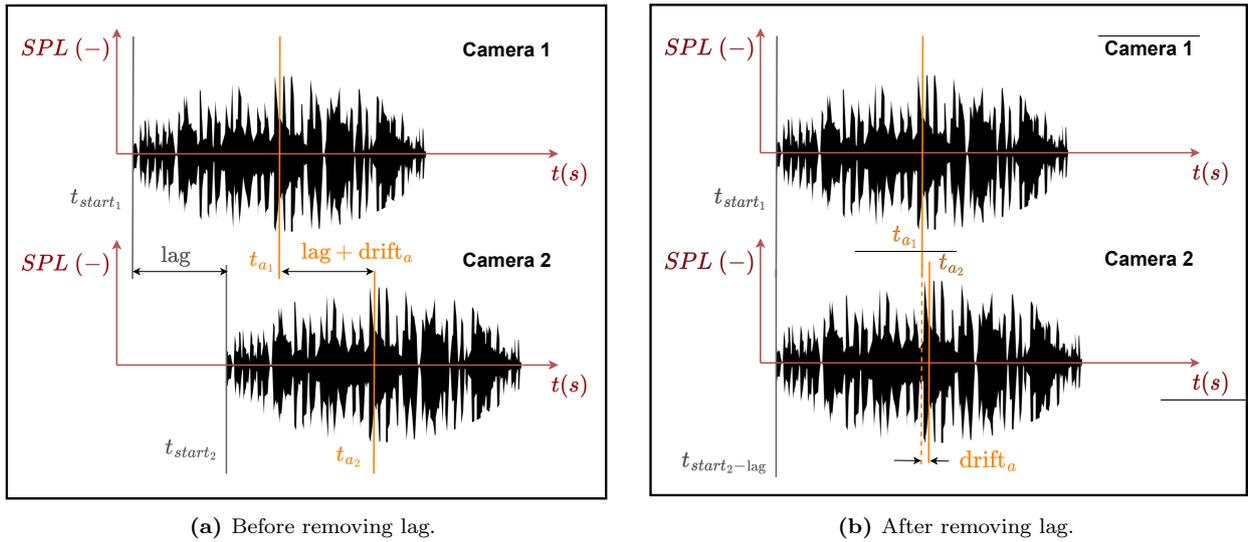


Figure 7.5: Audio-based synchronisation of GoPro recordings. (a) Original, unsynchronised signals. (b) After lag correction.

The drift at a later moment in time a is then

$$\text{drift}_a = \left| \frac{\Delta t_a - \text{lag}}{T_{\text{frame}}} \right|.$$

After correcting for lag, the audio waveforms align and the corresponding video frames can be matched reliably. An example from the final experiment is shown in Fig. 7.6, where seven claps are clearly aligned. Here, the lag has been corrected for. The output is a file of matching frame indices for the entire recording.

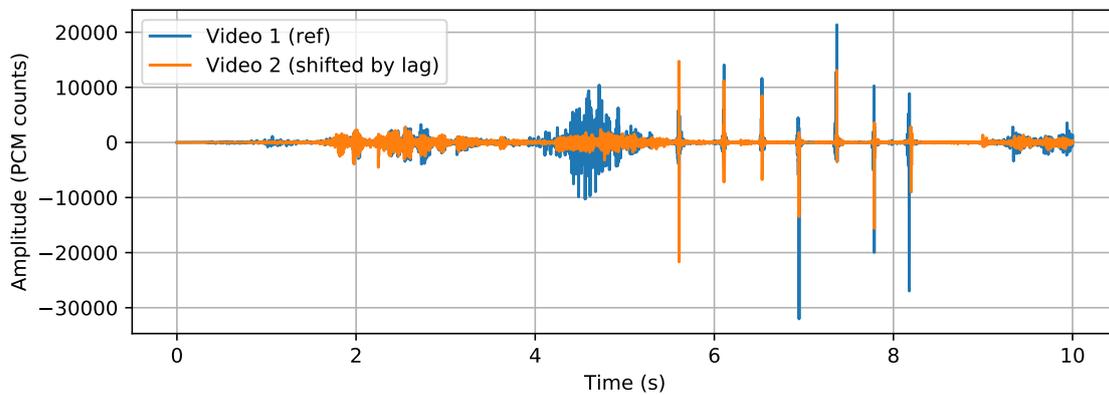


Figure 7.6: Synchronised sound waves from the final experiment in Bangor-Erris.

Drift analysis

To determine whether drift requires correction in addition to lag, a 15-minute video was split into 90-second segments. Within each segment, the relative offset between the two audio tracks was computed. The resulting drift curve is shown in Fig. 7.7. Over 15 minutes, the maximum drift remained below 0.1 frames—equivalent to only a few milliseconds—well below the threshold at which frame mismatches would occur (± 0.5 frame).

Based on these results, drift is assumed negligible for the expected experiment duration. Synchronisation is therefore performed by correcting only for lag.

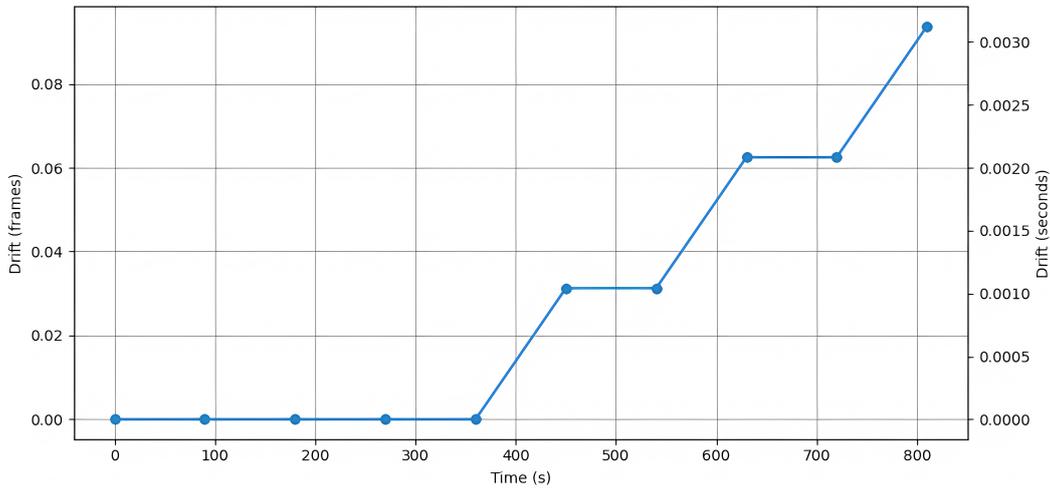


Figure 7.7: Measured drift between two GoPro cameras over a 15-minute period. Values remain far below the frame-mismatch threshold.

7.4.2 KCU and UWB synchronisation

The UWB modules are initialised over SSH from a laptop whose system clock is synchronised to Coordinated Universal Time (UTC) via the Network Time Protocol (NTP). The UWB timestamps therefore inherit the UTC timing of the Raspberry Pi controlling the ranging sequence.

The Pitot tube is hard-wired to the KCU, so Pitot and KCU data share the same internal clock. The KCU communicates continuously with the ground station, which is synchronised to an external time server. As a result, both the UWB and KCU data streams are referenced to UTC throughout flight.

7.4.3 Photogrammetry and UWB synchronisation

After synchronising the two cameras with each other, they must be linked to the UTC-based sensors. Because the GoPro does not provide UTC timestamps, a synchronisation flash is used.

At the moment UWB ranging starts, the Raspberry Pi activates an LED:

```
1 import RPi.GPIO as GPIO
2 GPIO.output(pin, GPIO.HIGH)
```

The LED sits inside the sensor box and is clearly visible to the GoPro. The brightness transition (dark → bright → dark) creates a sharp temporal marker used to map video frames to UTC time. This sequence is virtually depicted in Fig. 7.8.

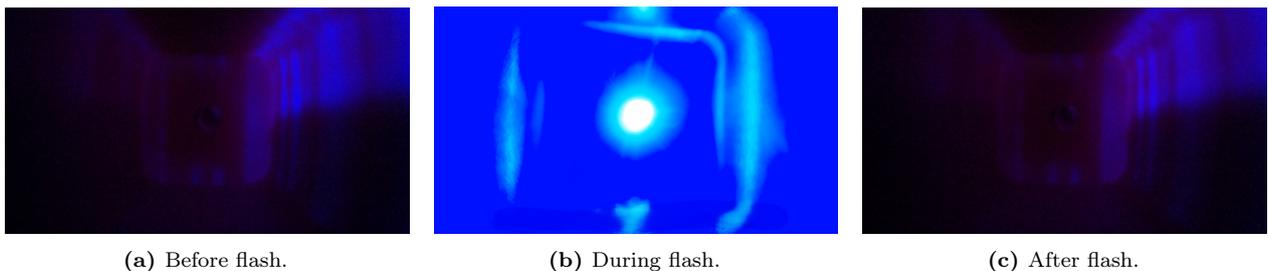


Figure 7.8: Flash sequence used to map GoPro frames to UTC.

The second camera is subsequently shifted by the same offset once the first is aligned.

Flash detection method

Flash onset is detected through a brightness-slope criterion. Let B_i denote mean brightness in frame i over a predefined ROI. Define

$$\Delta B_i = B_i - B_{i-1}.$$

The flash frame i_f is the first frame satisfying

$$\Delta B_i > \kappa \cdot \overline{\Delta B},$$

with $\kappa = 5$ and $\overline{\Delta B}$ the average preceding slope. The result is shown in Fig. 7.9.

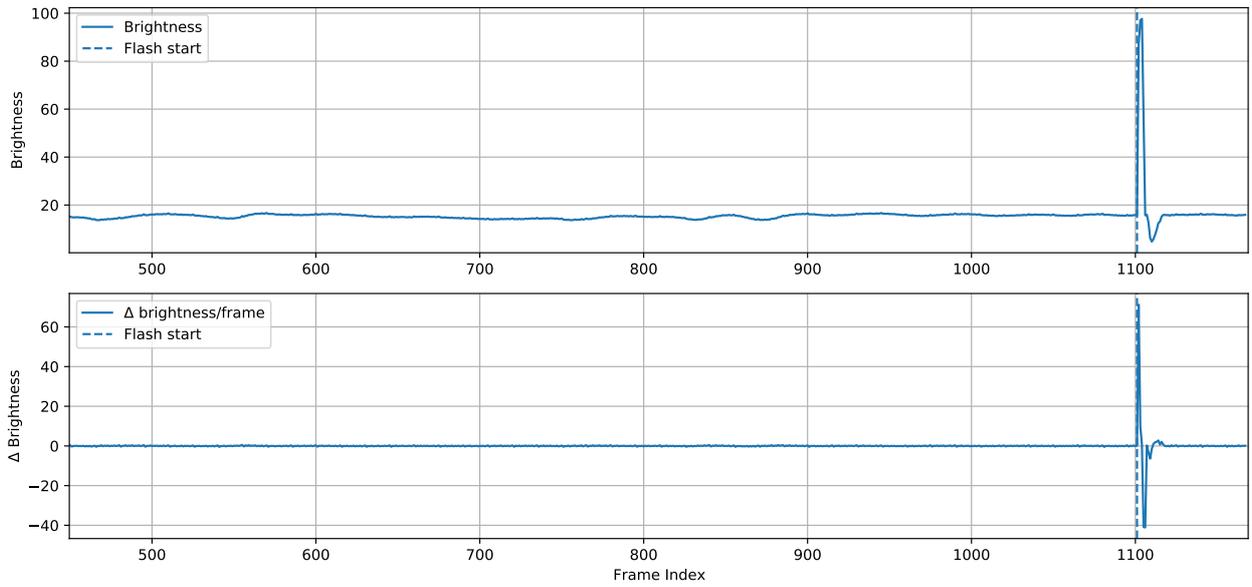


Figure 7.9: Flash onset detection via brightness slope. The vertical line marks the synchronisation frame.

Time coupling

Once i_f is known, each GoPro frame is assigned a UTC timestamp

$$t_i = \frac{i - i_f}{f_{\text{cam}}},$$

using the nominal frame rate $f_{\text{cam}} = 29.97$ Hz, obtained by:

```
1 cap = cv2.VideoCapture(video_path)
2 cap.get(cv2.CAP_PROP_FPS)
```

It is assumed that the GoPro maintains this nominal rate and that drift relative to UTC-based sensors remains negligible for the experiment duration.

LED-UWB temporal offset

A small delay exists between the actual LED flash and the first UWB UTC entry. This was quantified by filming an external UTC display during the flash. The measured offset is shown in Fig. 7.10.

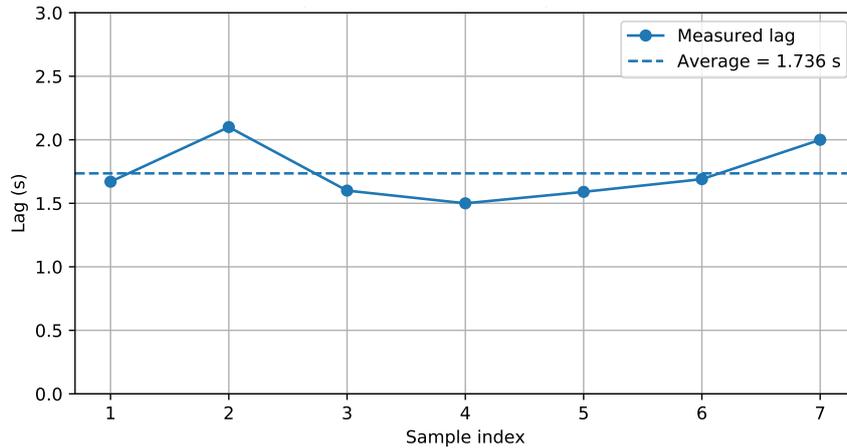


Figure 7.10: Lag between LED flash (true UTC) and the first UWB UTC log. Mean offset is 1.736 s.

The average lag (1.736 s) is added to the UWB timestamps. Although the offset varies between 1.5 – 2.1 s, this uncertainty is acceptable: UWB and photogrammetry are aligned precisely via the flash, and KCU/Pitot data vary slowly enough that sub-second discrepancies do not affect the conclusions. For the present work, all sensors are therefore treated as synchronised.

7.5 Measurement architecture

Figure 7.11 provides a schematic overview of the complete experimental setup, including sensors, data flows, power supplies, and synchronisation signals. The text below describes the overview.

The experiment begins with the ground station (GS), which communicates commands to the lithium-ion battery powered KCU. The KCU translates these commands into depower and steering inputs, u_{dp} and u_s , which, together with the prevailing environmental conditions, drive the kite’s motion and deformation. The apparent wind speed V_a is measured by a Pitot tube. The front-line angle of attack α_{fl} and sideslip angle β_s are measured by the angular vanes. This sensor subsystem is wired into the KCU, which also provides time-stamped output data. The KCU simultaneously logs all sort of data, from which the position of the kite in spherical coordinates (ϕ, β, r) and tether force F_T are saved for this research.

Wing shape is measured by the photogrammetry subsystem, consisting of two sound-synchronised cameras. These cameras are linked to the universal timeline through a synchronisation flash generated by an LED connected to the Raspberry Pi 3A+. The image data are processed through the photogrammetry pipeline to obtain time-resolved 3D marker coordinates $(x_{i,w}, y_{i,w}, z_{i,w})$ and labelled datasets, so a wireframe model of the wing can be reconstructed for each frame.

In parallel, the UWB subsystem measures the instantaneous span b . A powerbank supplies a Raspberry Pi 3A+, which powers the UWB shield running Pozyx ranging code. The shield transmits wireless commands to wingtip-mounted UWB anchors, which perform active ranging and send the span measurement back to the shield. The span data are stored locally on the Raspberry Pi SD card.

In addition to the onboard measurements, the ground station also records a subset of relevant flight data. The wing is equipped with a Kitepower-integrated sensor module that provides measurements such as kite position, tether force, and various auxiliary signals. However, not all available channels are used for post-processing. Only the kite’s spherical coordinates (ϕ, β, r) and the tether force F_T were retained, as these were considered the most relevant for the analysis in this study.

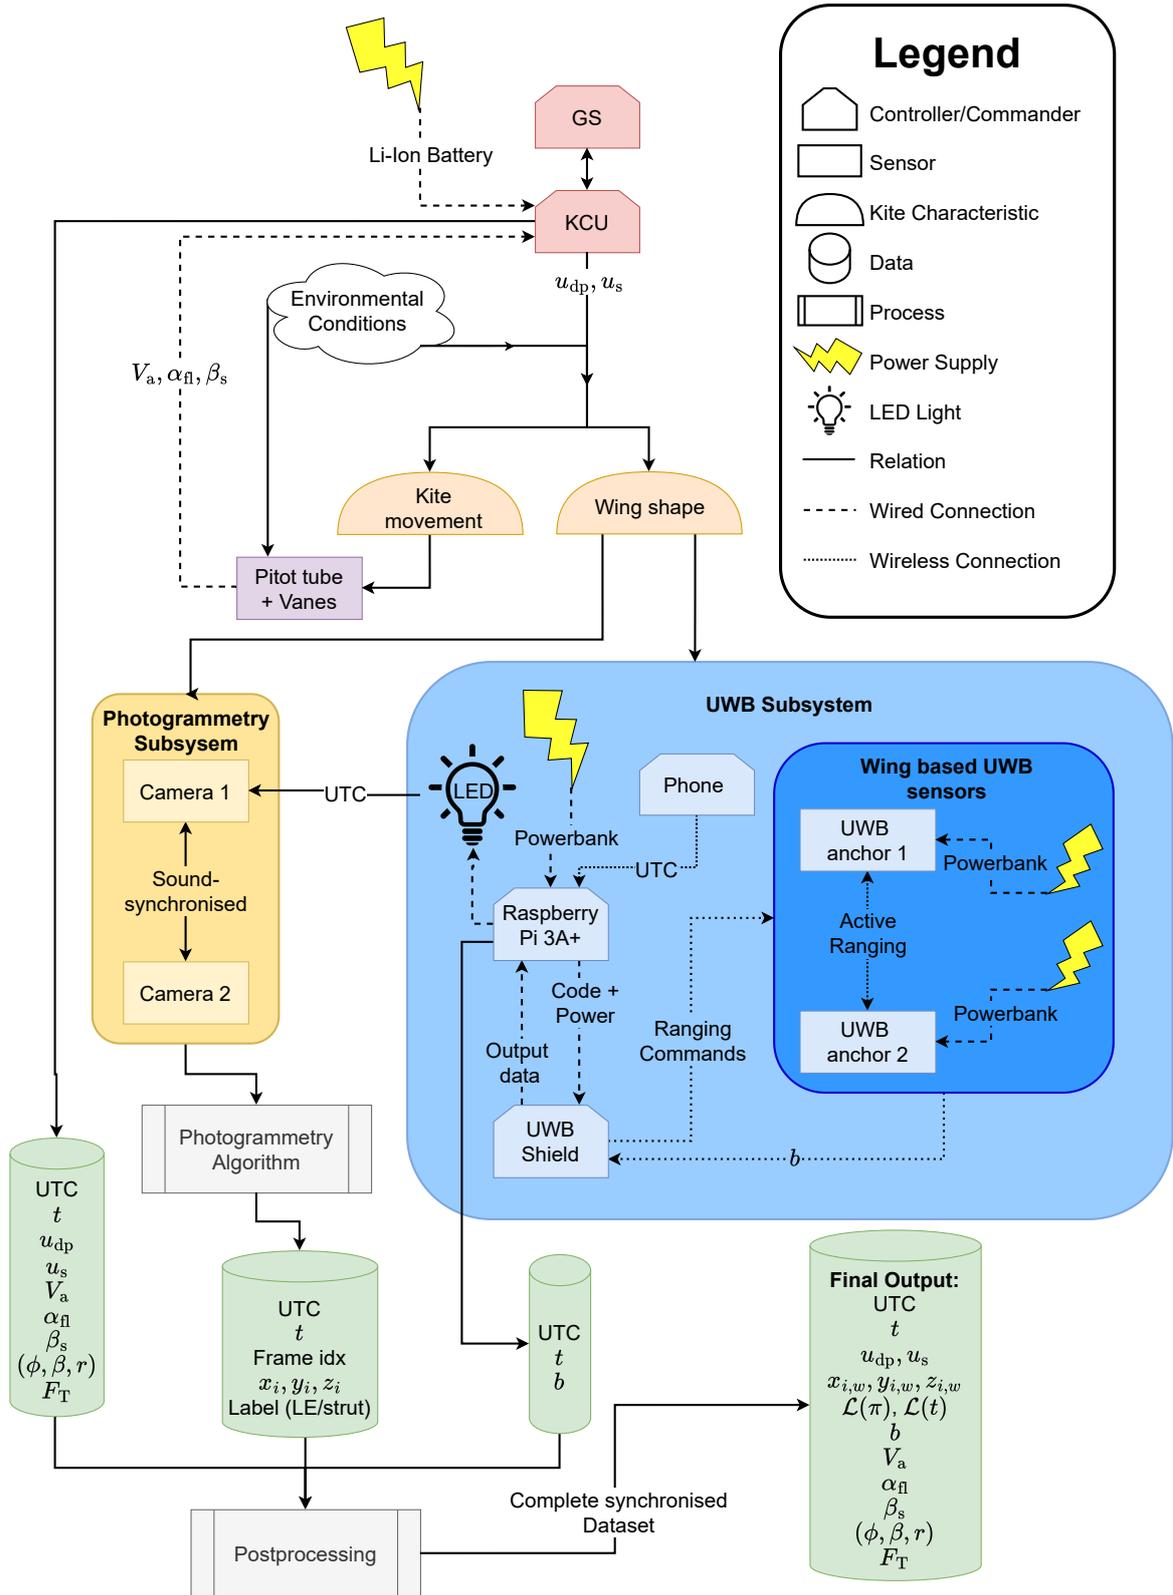


Figure 7.11: Schematic overview of the experimental setup. The figure illustrates the full measurement architecture, including sensors, data flows, and synchronisation links. The GS issues commands to the KCU, which actuates the kite and interfaces with the Pitot tube and angular vanes. Wing shape is measured independently by the photogrammetry subsystem and the UWB subsystem. All outputs are postprocessed and synchronised to UTC to form the final dataset.

After postprocessing and synchronisation, the final dataset is described in Table 7.4. Together, these outputs constitute a complete, time-resolved description of the kite's state during flight. The post-processing block, which generates a point cloud in the wing coordinate frame and subsequently

constructs a wireframe model, is further specified in Chapter 8.

Table 7.4: Model outputs, all subsystems synchronised and combined

Output	Symbol
UTC time	UTC
Timestamp	t
Depower input	u_{dp}
Steering input	u_s
Marker coordinates	$x_{i,w}, y_{i,w}, z_{i,w}$
Wireframe model	$\mathcal{L}(\pi), \mathcal{L}(t)$
Span	b
Apparent wind speed	V_a
Front line angle of attack	α_{fl}
Sideslip angle	β_s
Kite position	(ϕ, β, r)
Tether force	F_T

Post-processing is required for all sensor systems defined in the measurement setup, with the exception of the data provided directly by the KCU. The KCU is a well-integrated subsystem developed and maintained by Kitepower B.V., and analysing, filtering and calibrating its data is beyond the scope of this research. Consequently, KCU signals are used as delivered and are not modified further in this work.

For all other data streams, post-processing serves three main purposes:

- (I) Correcting systematic measurement bias through **linear calibration**.
- (II) **Handling missing samples** so that all time series can be analysed on a common timeline.
- (III) **Reducing high-frequency noise** without introducing artificial time lag between signals.

These generic operations are introduced in Section 8.1; later sections specify how they are applied to each sensor and describe additional, photogrammetry-specific geometric corrections. In particular, the transformation to a wing-fixed reference frame, wing shape reconstruction, and stereoscopic calibration corrections are grouped under photogrammetry-specific post-processing at the end of this chapter (Section 8.2).

8.1 General post-processing operations

This section summarises the core post-processing operations that are reused across multiple data streams. Where possible, the mathematical definitions are given in a standard form and summarised in a compact table later in this chapter.

8.1.1 (I) Linear calibration of raw signals

Most one-dimensional sensor outputs are calibrated using a linear relationship between the raw digital reading x_{raw} and the physical quantity $x_{\text{calibrated}}$ of interest:

$$x_{\text{calibrated}} = a x_{\text{raw}} + b, \quad (8.1)$$

where a is a scale factor and b is an offset. The constants (a, b) have been determined experimentally for the UWB modules, the pitot tube and the angular vanes and are reported in the corresponding chapters (see Chapter 5 and Chapter 6). With (a, b) known, (8.1) is applied pointwise to the raw time series to obtain calibrated values in physical units.

For vector-valued quantities, such as 3D coordinates from photogrammetry, the term “calibration” is used in a broader sense and refers to the individual and stereo camera calibration and geometric correction procedure. Those steps are described in detail in Chapter 4 and the photogrammetry-specific post-processing sections later in this chapter, and are not repeated here.

8.1.2 (II) Handling missing samples

Short gaps occur in several time series, for example when the UWB link momentarily drops out or when a photogrammetry marker is not detected in a particular frame. To obtain continuous signals on the common timeline, missing samples are filled by linear interpolation, provided that the gap is sufficiently short.

Let $\mathbf{p}(t)$ denote a scalar or vector-valued measurement at time t (for instance, a 3D marker coordinate). Suppose that a sample at time t is missing, while valid measurements exist at times $t_1 < t < t_2$. Then the missing value is reconstructed as

$$\mathbf{p}(t) = \mathbf{p}(t_1) + \frac{t - t_1}{t_2 - t_1} (\mathbf{p}(t_2) - \mathbf{p}(t_1)).$$

In practice, interpolation is only applied when the gap length ($t_2 - t_1$) is below a signal-specific threshold such that the underlying dynamics can be assumed approximately linear over the gap. Longer gaps are either left as missing values or excluded from subsequent analysis, depending on the application.

For the UWB span signal, only isolated missing samples or short runs of at most a few samples are interpolated. For photogrammetry, linear interpolation is applied only to markers whose 3D position changes very little over several consecutive frames, indicating that they correspond to the same physical marker. Two detections are treated as the same marker only if their separation stays below $r_{\text{match}} \leq 0.25$ m over the considered frames.

This threshold is smaller than the minimum spacing between neighbouring markers, which prevents mixing of adjacent points and ensures consistent marker identification even though the marker indices themselves are not temporally consistent. In this way, interpolation improves temporal continuity of the reconstructed geometry without introducing spurious connections between different markers.

8.1.3 (III) Reducing high-frequency noise

Even after calibration and interpolation, most signals exhibit high-frequency jitter around a smoother underlying trend. To suppress this noise while preserving the physically relevant dynamics, an exponential moving average (EMA) is used as a low-pass filter [136]. Because all post-processing is performed offline, one-dimensional time series can be filtered in a forward-backward (zero-phase) fashion to avoid any effective time lag. For photogrammetry, a spatially constrained EMA is applied per 3D point, as detailed below. This filtering method ensures the data can still be used in synchronisation with the KCU output, without lag.

Forward EMA filtering

For a discrete-time input signal $x[n]$ with sampling period Δt , a first-order EMA produces a filtered output $y^f[n]$ according to the recursion, where superscript f indicates the forward-filtered signal:

$$y^f[n] = \alpha y^f[n - 1] + (1 - \alpha) x[n], \quad 0 \leq \alpha < 1, \quad (8.2)$$

where α is the smoothing factor. Large values of α place more weight on the previous output and therefore yield stronger smoothing but slower response to changes; smaller values of α result in a more responsive but less smooth signal. The filter is initialised by setting $y^f[0] = x[0]$ for each processed time series.

Taking the z -transform of (8.2) yields

$$Y(z) = \alpha z^{-1} Y(z) + (1 - \alpha) X(z),$$

and thus the discrete-time transfer function is

$$H(z) = \frac{Y(z)}{X(z)} = \frac{1 - \alpha}{1 - \alpha z^{-1}}.$$

The frequency response is obtained by evaluating the transfer function on the unit circle, $z = e^{j\omega}$:

$$H(e^{j\omega}) = H(z)|_{z=e^{j\omega}} = \frac{1 - \alpha}{1 - \alpha e^{-j\omega}}.$$

This corresponds to a first-order IIR low-pass magnitude response and a non-zero phase response (i.e., a small effective delay) for a single forward pass.

Backward EMA filtering and zero-phase construction

To remove this delay and keep signals aligned in time, the EMA is applied twice: once in the forward time direction and once in reverse. Let $y^f[n]$ denote the output of the forward EMA pass applied to $x[n]$ using (8.2). A second EMA is then applied to the time-reversed sequence $y^f[N - n]$, yielding a backward-filtered sequence $y^b[n]$. Reversing the result again gives the final zero-phase output $\tilde{y}[n]$. In the frequency domain, the combined forward-backward operation has transfer function

$$H_0(e^{j\omega}) = H(e^{j\omega}) H(e^{-j\omega}) = |H(e^{j\omega})|^2,$$

which has zero phase for all frequencies and an enhanced low-pass magnitude response compared to a single EMA. As a result, high-frequency jitter is strongly attenuated without shifting peaks or transients in time.

Spatially constrained EMA for photogrammetry markers

For photogrammetry, a simple sample-wise EMA would blur unrelated points if the correspondence between frames were not enforced. Therefore, a spatially constrained EMA is applied to the 3D wing-frame coordinates, combining temporal smoothing with nearest-neighbour matching.

Let $\{\mathbf{p}_t^k\}$ denote the set of wing-frame points in frame t , and let $\{\tilde{\mathbf{p}}_{t-1}^j\}$ be the previously filtered points in frame $t - 1$. For each current point \mathbf{p}_t^k , the closest predecessor is determined:

$$j^* = \arg \min_j \|\mathbf{p}_t^k - \tilde{\mathbf{p}}_{t-1}^j\|_2.$$

If this minimum distance is smaller than a matching radius r_{match} , the two points are interpreted as the same physical marker and an EMA is applied in 3D:

$$\tilde{\mathbf{p}}_t^k = \begin{cases} \alpha \tilde{\mathbf{p}}_{t-1}^{j^*} + (1 - \alpha) \mathbf{p}_t^k, & \|\mathbf{p}_t^k - \tilde{\mathbf{p}}_{t-1}^{j^*}\|_2 \leq r_{\text{match}}, \\ \mathbf{p}_t^k, & \text{otherwise.} \end{cases}$$

Points without a suitable predecessor (for example, newly appearing markers or strong outliers) are left unchanged. In this work, α is chosen such that frame-to-frame jitter is reduced while preserving the physically relevant motion of the wing, and $r_{\text{match}} = 0.25$ m, set smaller than the minimum spacing between neighbouring markers to avoid incorrect associations.

The resulting sequence $\{\tilde{\mathbf{p}}_t^k\}$ is a clean, temporally consistent set of wing-frame coordinates suitable for line fitting on the struts and ordering along the leading edge.

Illustrative effect of zero-phase EMA

The effect of the zero-phase EMA on a typical noisy time series is illustrated in Fig. 8.1. The example shows UWB distance fluctuations around the mean span during a stationary test: the raw signal exhibits frame-to-frame variations of several centimetres, while the filtered signal closely follows the slowly varying trend without any visible time lag relative to the raw data.

8.1.4 Overview of applied post-processing per signal

The generic operations defined in Section 8.1.1 and Section 8.1.3 are not applied uniformly to all data streams. Table 8.1 summarises, for each signal, which post-processing steps are used. Details and numerical values (e.g. calibration constants) are given in the corresponding sections referenced in the table.

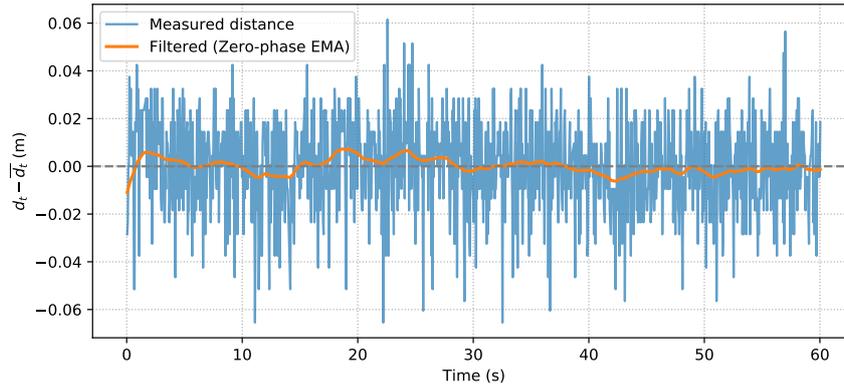


Figure 8.1: Example of zero-phase EMA filtering on UWB span data. The mean distance has been subtracted to show only fluctuations around the mean. The raw measurements (blue) exhibit high-frequency jitter, while the zero-phase EMA output (orange, $\alpha = 0.9$) provides a smoother trace without apparent time shift.

Table 8.1: Overview of post-processing operations applied to each data stream.

Data stream	(I) Calibration	(II) Interpolation of gaps	(III) Temporal filtering
Photogrammetry (markers)	Individual and stereo calibration (Chapter 4).	Linear interpolation of short gaps for reference markers and slowly moving markers.	Spatial EMA with nearest-neighbour matching, using radius $r_{\text{match}} = 0.25$ m and smoothing factor $\alpha = 0.95$.
UWB span	Linear calibration of span (Section 5.3: $a_{\text{UWB}} = 0.9992$, $b_{\text{UWB}} = 0.0793$).	Linear interpolation of isolated missing samples or short runs.	Zero-phase EMA, $\alpha = 0.95$.
Pitot tube	Linear calibration of V_∞ (Section 6.1.2: $a_{\text{ws}} = 1.064$, $b_{\text{ws}} = -0.415$).	Not required (timestamps always paired with V_∞).	Zero-phase EMA, $\alpha = 0.9$.
Angular vanes	Linear calibration of in-flow angles (Section 6.2.1: $a_{\text{av}} = 0.9195$, $b_{\text{av}} = 0.4202$).	Not required (regular, complete sampling).	Zero-phase EMA, $\alpha = 0.9$.
KCU	Internal KCU processing.	None.	None.

8.2 Photogrammetry-specific post-processing

Beyond the generic operations of calibration, interpolation and temporal filtering, the photogrammetry pipeline requires several additional geometric post-processing steps. These operate directly on the triangulated 3D marker coordinates and are applied in the following order:

- (i) **Correction** of the stereo geometry for **bending and twist** of the camera rig.
- (ii) **Transformation** of all reconstructed points to a **wing-fixed reference frame**.
- (iii) **Removal** of residual **outliers** in the 3D point clouds.
- (iv) **Wing shape reconstruction** from the cleaned and transformed marker set.

The first step modifies the effective stereo calibration; the remaining steps are applied to the resulting 3D marker coordinates on a frame-by-frame basis after interpolation and temporal filtering.

8.2.1 (i) Correction for camera-rig bending and twist

The initial photogrammetric reconstructions showed strong inconsistencies when compared to the reference distances obtained from the UWB sensors. The reconstructed span of the kite did not correspond to the UWB-measured span, and the discrepancies were both systematic and large. After multiple verification tests, the UWB measurements were found to be highly reliable, leading to the conclusion that the problem originated within the photogrammetry pipeline itself. This conclusion was supported by additional observations. First, the reconstructed strut lengths—which should remain constant due to the rigid carbon reinforcements—varied significantly between frames. Second, the epipolar constraints had to be relaxed to unrealistically high thresholds (up to $|\delta_y| > 50$ pixels) for corresponding markers to be matched, which is inconsistent with a well-calibrated stereo system.

Furthermore, the estimated kite geometry varied between test conditions. During flight, the reconstructed kite appeared smaller and closer to the camera rig than expected, while during the static “hanging-kite” tests (see Chapter 10), the reconstructed span was overestimated. Together with the verified stability of the UWB measurements, these findings indicate that the glass-fibre support bar was subject to slight elastic deformation under load, altering the relative orientation of the cameras. Since the experiments could not be repeated within the duration of the thesis, a post-processing correction was required.

The Γ -Bar: Bending and Twist

The glass-fibre Γ -bar with the 3D-clamps was designed to be held under moderate bridle tension, assuming negligible deformation during operation. This was supposed to be ensured by tactical placement of the bar, such that loads were minimized. However, variations in bridle tension, orientation and aerodynamic loading caused the bar to experience small bending and twisting deformations. Although these angular deviations were small ($< 2^\circ$), their effect on stereo triangulation was substantial due to the long baseline and large object distance. Consequently, even minor orientation changes could introduce scale errors or vertical misalignment between images.

To quantify these effects and compensate for them, the influence of small relative rotations between the cameras was analysed for each degree of freedom: yaw, roll, and pitch. Two sets of angular quantities are defined throughout this section. The first, denoted $\Delta\mu_{\text{bend}}$ ($\mu \in \{\psi, \phi, \theta\}$), describes the actual misalignment of the cameras during capture, caused by elastic deformation of the Γ -bar. The second, $\Delta\mu_{\text{corr}}$, represents the angular correction applied in post-processing to restore the nominal calibration. By definition, the correction is opposite in sign to the bending,

$$\Delta\mu_{\text{corr}} = -\Delta\mu_{\text{bend}},$$

so that the applied rotation compensates for the deformation-induced misalignment. The angles depicted in the illustrative images therefore show both $\Delta\mu_{\text{bend}}$ and $\Delta\mu_{\text{corr}}$, clarifying that they are equal in magnitude but opposite in sign.

Yaw Effect (ψ): yaw misalignment was found to have the largest influence on the reconstructed geometry, which is expected given the camera placement along the kite span. Figure 8.2 illustrates how relative yaw between the cameras affects horizontal disparity and perceived depth.

In the nominal configuration (a), corresponding to the calibrated stereo geometry, the camera orientations match the rotation matrix \mathbf{R} (i.e., $\Delta\psi_{\text{bend}} = \Delta\psi_{\text{corr}} = 0$). The triangulated points yield the correct 3D geometry, with mean vertical disparity $|\delta_y| \approx 0$ and accurate horizontal disparity δ_x .

When the cameras converge due to bar bending (b), the misalignment is $\Delta\psi_{\text{bend}} < 0$. The post-processing applies the opposite correction, $\Delta\psi_{\text{corr}} > 0$. Convergence reduces the horizontal disparity δ_x , leading to an overestimation of depth and an apparent enlargement of the reconstructed kite.

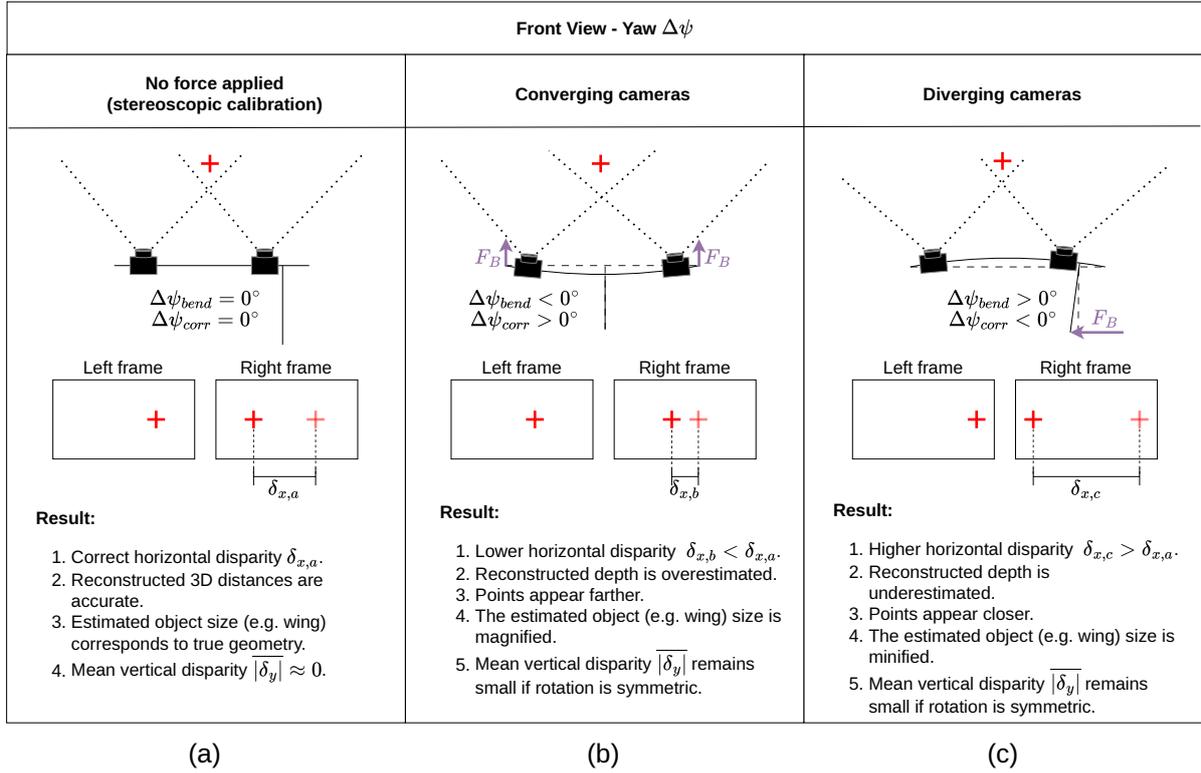


Figure 8.2: Effect of yaw change due to bending of the Γ -bar. (a) Nominal stereoscopic calibration ($\Delta\psi_{bend} = \Delta\psi_{corr} = 0$); (b) converging cameras ($\Delta\psi_{bend} < 0$, $\Delta\psi_{corr} > 0$) resulting in overestimated depth and enlarged scale; and (c) diverging cameras ($\Delta\psi_{bend} > 0$, $\Delta\psi_{corr} < 0$) resulting in underestimated depth and reduced scale. The correction angles $\Delta\psi_{corr}$ are applied in post-processing with opposite sign to the deformation-induced misalignment.

Conversely, for diverging cameras (c) we have $\Delta\psi_{bend} > 0$ and $\Delta\psi_{corr} < 0$, which increases δ_x , underestimates depth, and yields a smaller reconstructed span.

Because yaw is a symmetric rotation around the vertical axis, the mean vertical disparity remains small; thus, yaw misalignment primarily results in an incorrect reconstruction scale and camera-to-object distance.

Roll Effect (ϕ): roll deviations were also observed, particularly during turning manoeuvres when asymmetric loads on the leading-edge bridles caused the bar to twist slightly about its longitudinal axis. Similar bending of the bar could also occur from drag forces at higher apparent wind speeds V_a , often resulting in a negative roll misalignment in flight (i.e., $\Delta\phi_{bend} < 0$). Figure 8.3 shows the qualitative effect of this rotation on the projected images.

In the calibrated case (a), both cameras are aligned such that $|\overline{\delta_y}| \approx 0$, and corresponding points lie on the same epipolar lines ($\Delta\phi_{bend} = \Delta\phi_{corr} = 0$). When the right camera is rotated clockwise by bending (b), $\Delta\phi_{bend} < 0$, and the correction is $\Delta\phi_{corr} > 0$. The image plane rotates around the optical axis: points on the left side of the frame (relative to the principal point (c_x, c_y)) move upward, while points on the right move downward, introducing opposite-signed vertical disparities across the image. The inverse effect occurs for a counter-clockwise misalignment (c), where $\Delta\phi_{bend} > 0$ and $\Delta\phi_{corr} < 0$.

These roll-induced disparities cause an in-plane shear distortion in the reconstructed geometry. The mean vertical disparity $|\overline{\delta_y}|$ increases approximately linearly with $|\Delta\phi_{bend}|$ (and equivalently with $|\Delta\phi_{corr}|$), even though horizontal disparity remains largely unaffected.

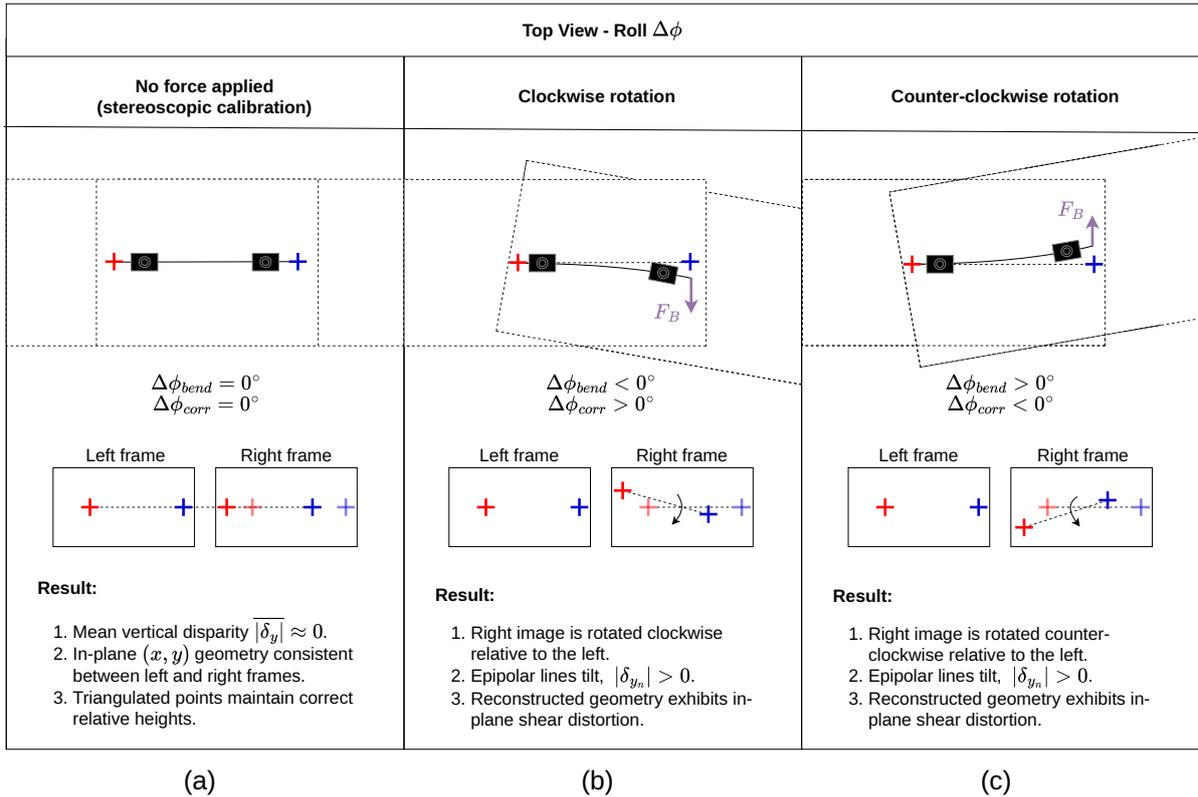


Figure 8.3: Effect of roll change caused by asymmetric bridle loading. (a) Nominal calibration ($\Delta\phi_{bend} = \Delta\phi_{corr} = 0$); (b) clockwise twist of the bar ($\Delta\phi_{bend} < 0$, $\Delta\phi_{corr} > 0$) producing opposite vertical disparities on either side of the frame; and (c) counter-clockwise twist ($\Delta\phi_{bend} > 0$, $\Delta\phi_{corr} < 0$) producing the inverse effect.

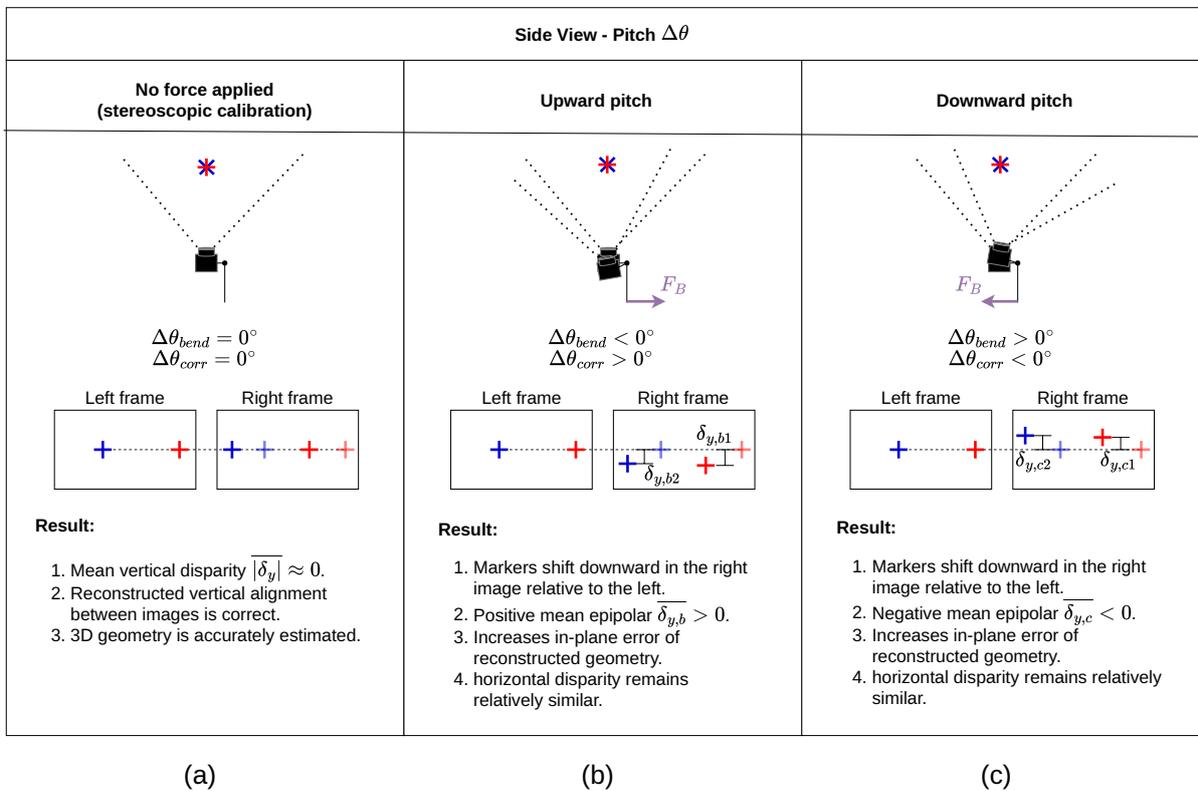


Figure 8.4: Effect of pitch change due to torsional deformation of the Γ -bar. (a) Nominal calibration ($\Delta\theta_{bend} = \Delta\theta_{corr} = 0$); (b) upward pitch ($\Delta\theta_{bend} < 0$, $\Delta\theta_{corr} > 0$) causing uniform downward marker shift and positive vertical disparity; and (c) downward pitch ($\Delta\theta_{bend} > 0$, $\Delta\theta_{corr} < 0$) causing uniform upward shift and negative vertical disparity.

Pitch Effect (θ): pitch misalignment, primarily caused by twisting of the Γ -bar, was generally smaller than yaw or roll deviations. Even so, its effect on the reconstructed geometry is distinct and measurable, as shown in Fig. 8.4.

In the nominal configuration (a), corresponding markers appear at the same height in both frames ($\Delta\theta_{\text{bend}} = \Delta\theta_{\text{corr}} = 0$), and the vertical disparity is approximately zero. When the right camera pitches upward by bending (b), $\Delta\theta_{\text{bend}} < 0$ and the correction is $\Delta\theta_{\text{corr}} > 0$. All corresponding markers in the right frame shift downward by approximately the same amount, resulting in a uniform positive vertical disparity and an increase in the mean epipolar error. Conversely, a downward pitch misalignment (c) has $\Delta\theta_{\text{bend}} > 0$ and $\Delta\theta_{\text{corr}} < 0$, shifting markers upward and producing a negative vertical disparity.

Since pitch primarily introduces a uniform vertical offset across the image, the horizontal disparity distribution and overall scale remain largely unaffected. The result is an in-plane error dominated by vertical alignment differences rather than depth misestimation.

Overview of effects

The analysis above demonstrates that each of the three small angular deviations between the stereo cameras—yaw, roll, and pitch—affects the reconstructed geometry in a distinct manner. Table 8.2 summarizes the dominant disparity components, the corresponding geometric errors, and their qualitative influence on the reconstructed 3D structure.

Table 8.2: Summary of the qualitative influence of small angular misalignments on disparity and reconstructed geometry.

Angle	Primary Disparity Affected	Geometric Effect	Qualitative Result
Yaw ($\Delta\psi$)	Horizontal disparity (δ_x)	Depth and scale error	Over- or underestimation of span; minimal epipolar error
Roll ($\Delta\phi$)	Vertical disparity gradient (δ_y varies with x)	In-plane shear distortion	Opposite vertical shift on either side of image; increased $ \delta_y $
Pitch ($\Delta\theta$)	Uniform vertical disparity (δ_y constant)	Vertical alignment offset	Uniform upward/downward marker shift; increased $ \delta_y $

Stereoscopic calibration correction

It is assumed that, under small angular deviations ($< 2^\circ$), the translation vector \mathbf{t} remains constant. The translational change caused by bar deformation is expected to be negligible compared to the effect of small orientation errors on depth reconstruction, which was confirmed through verification tests; the translation was much less sensitive than the rotation. This simplification is also computationally advantageous, as including \mathbf{t} would increase the number of variables to be optimized by three and significantly raise the computational complexity. As a result, only yaw, roll, and pitch corrections were determined and integrated into the rotational matrix \mathbf{R} . The assumption can later be validated by evaluating the residual error after applying the angular corrections (see Section 8.2.1).

To correct for the bending and twisting of the bar, an algorithm was developed that determines the small angular corrections between the cameras. These corrections are expressed as $\Delta\psi_{\text{corr}}$, $\Delta\phi_{\text{corr}}$, and $\Delta\theta_{\text{corr}}$, corresponding to yaw, roll, and pitch respectively. The aim is to find the combination of these correction angles that restores the correct stereo geometry.

Search grid: to achieve this, a three-stage grid search is performed over these three angular correction parameters. The algorithm starts with a broad search and iteratively refines it:

1. **Coarse search:** $\Delta\psi_{\text{corr}}$, $\Delta\phi_{\text{corr}}$, and $\Delta\theta_{\text{corr}}$ are varied with a step size of 0.5° .
2. **Medium search:** A smaller window around the best coarse result is explored with a step size of 0.05° .
3. **Fine search:** A final refinement is performed around the best medium-stage result with a step size of 0.005° .

In this procedure, only one set of manual correspondences is required. This procedure requires manual selection of corresponding markers on the uncalibrated left and right frames. The use of unrectified images is essential, since the calibration itself is being corrected: rectified images would already impose an (incorrect) geometry based on the initial, misaligned rotation matrices. Using raw images ensures that the correction directly reflects the true geometric deviation.

Mathematical implementation: mathematically, the correction is applied to the rotational matrix \mathbf{R} obtained from the original stereoscopic calibration. A small rotation $\Delta\mathbf{R}$ is constructed using the three incremental correction angles:

$$\Delta\mathbf{R} = \mathbf{R}_y(\Delta\psi_{\text{corr}}) \mathbf{R}_x(\Delta\theta_{\text{corr}}) \mathbf{R}_z(\Delta\phi_{\text{corr}}),$$

where \mathbf{R}_x , \mathbf{R}_y , and \mathbf{R}_z represent elementary rotation matrices around the x -, y -, and z -axes respectively. The corrected rotation matrix \mathbf{R}' used for rectification then becomes:

$$\mathbf{R}' = \Delta\mathbf{R}\mathbf{R}.$$

Each elementary rotation matrix is defined as:

$$\mathbf{R}_x(\Delta\theta_{\text{corr}}) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\Delta\theta_{\text{corr}}) & -\sin(\Delta\theta_{\text{corr}}) \\ 0 & \sin(\Delta\theta_{\text{corr}}) & \cos(\Delta\theta_{\text{corr}}) \end{bmatrix},$$

$$\mathbf{R}_y(\Delta\psi_{\text{corr}}) = \begin{bmatrix} \cos(\Delta\psi_{\text{corr}}) & 0 & \sin(\Delta\psi_{\text{corr}}) \\ 0 & 1 & 0 \\ -\sin(\Delta\psi_{\text{corr}}) & 0 & \cos(\Delta\psi_{\text{corr}}) \end{bmatrix},$$

$$\mathbf{R}_z(\Delta\phi_{\text{corr}}) = \begin{bmatrix} \cos(\Delta\phi_{\text{corr}}) & -\sin(\Delta\phi_{\text{corr}}) & 0 \\ \sin(\Delta\phi_{\text{corr}}) & \cos(\Delta\phi_{\text{corr}}) & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

These matrices correspond to the right-handed rotations around the x , y , and z axes of the left camera frame, with positive angles defined according to the right-hand rule. The final correction $\mathbf{R}' = \Delta\mathbf{R}\mathbf{R}$ thus represents a small, compounded rotation applied to the right camera relative to the left camera, consistent with how the bar's bending or twisting would change the relative orientation of the two cameras in space.

Criteria: for each adjusted configuration, the algorithm performs stereo rectification and projects the manually selected marker centres into rectified space using \mathbf{R}' . The performance of each configuration is then quantified using two metrics:

1. **Span deviation:** the difference between the triangulated span (b) and the UWB-measured span. This metric is particularly sensitive to yaw correction errors ($\Delta\psi_{\text{corr}}$), since yaw primarily affects horizontal disparity and thus the perceived scale of the reconstructed kite.
2. **Mean vertical disparity:** the mean absolute vertical disparity $\overline{|\delta_y|}$ between the left and right rectified points. This metric is primarily sensitive to roll ($\Delta\phi_{\text{corr}}$) and pitch ($\Delta\theta_{\text{corr}}$) corrections, as these introduce vertical misalignments in the image plane. The absolute values are used because symmetric roll around the principal point may yield $\overline{\delta_y} \approx 0$, as positive and negative disparities can cancel each other out. For simplicity, $\overline{|\delta_y|}$ will be referred to as the epipolar error throughout this work, since both positive and negative disparities represent misalignment.

Both criteria are necessary. Minimizing only the span deviation would yield a correct scale but might leave substantial vertical misalignment (epipolar error), resulting in poor triangulation accuracy.

Conversely, minimizing only the vertical disparity could yield visually well-aligned images but still produce a distorted span and depth reconstruction. Using both requirements ensures all three angular corrections are computed accurately. Therefore, the optimal solution is determined through a combined criterion:

1. First, all solutions where the span deviation $|\Delta b|$ lies within ± 0.05 m of the UWB span are considered.
2. Among these, the one with the lowest mean vertical disparity $\overline{|\delta_y|}$ is selected.
3. If no solution meets the span tolerance, the smallest span violation is accepted, again minimizing the vertical disparity.

The result of this procedure is the set of optimal angular corrections $(\Delta\psi_{\text{corr}}^*, \Delta\phi_{\text{corr}}^*, \Delta\theta_{\text{corr}}^*)$, which can then be used to update the rotation matrix \mathbf{R}' for the reconstruction of that particular frame. This method effectively fuses geometric information from photogrammetry and distance measurements from the UWB system, ensuring that both scale and epipolar consistency are simultaneously satisfied. The optimization framework thus compensates for the small mechanical deformations of the camera rig while maintaining the integrity of the stereoscopic calibration.

Results

An example of the outcome of the optimization procedure is shown in Fig. 8.5, where both the mean vertical disparity and the span deviation relative to the UWB measurements are plotted for different combinations of yaw, roll, and pitch corrections. The results confirm the theoretical trends discussed earlier. The span deviation is primarily governed by the yaw correction $\Delta\psi_{\text{corr}}$, while the roll $\Delta\phi_{\text{corr}}$ and pitch $\Delta\theta_{\text{corr}}$ angles mainly affect the vertical alignment and thus the mean epipolar error. The presence of a consistent and unique optimum in each case supports the hypothesis that the observed geometric inconsistencies originated from small angular deviations caused by bending and twisting of the Γ -bar.

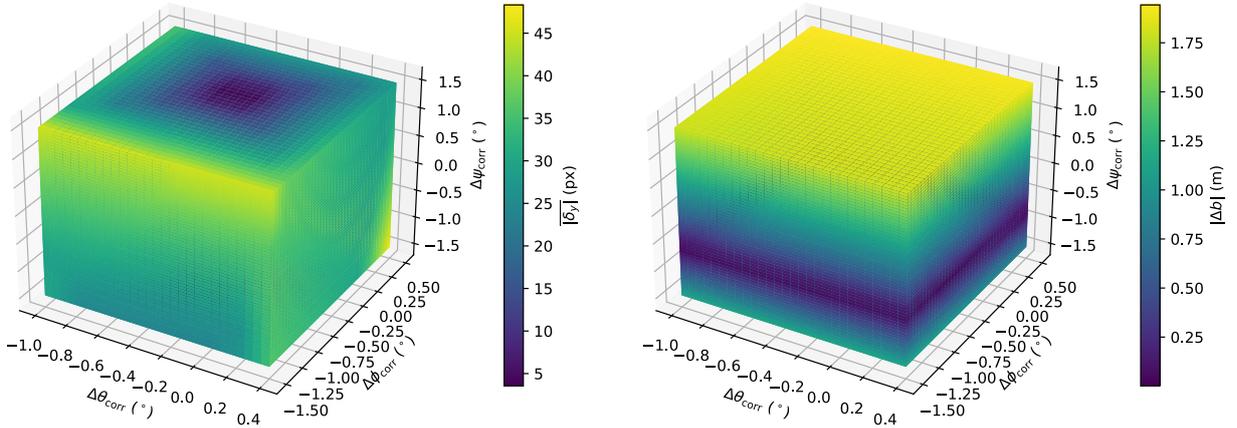


Figure 8.5: Three-dimensional parameter study showing the influence of small angular corrections on reconstruction quality. The colour scale indicates (left) the average vertical disparity $\overline{|\delta_y|}$ and (right) the span deviation Δb between photogrammetry and UWB measurements. The results confirm that span errors are predominantly driven by yaw correction ($\Delta\psi_{\text{corr}}$), while roll ($\Delta\phi_{\text{corr}}$) and pitch ($\Delta\theta_{\text{corr}}$) primarily influence vertical alignment.

A cross-section through the optimized yaw correction $\Delta\psi_{\text{corr}}^*$ is shown in Fig. 8.6. The left plot displays the mean vertical disparity as a function of $\Delta\phi_{\text{corr}}$ and $\Delta\theta_{\text{corr}}$, while the right plot shows the corresponding span deviation. The epipolar alignment varies strongly with roll and pitch, whereas yaw contributes only marginally; consistent with the theoretical discussion. Similarly, the span deviation remains almost constant with pitch but exhibits mild sensitivity to roll, which can be explained by small asymmetries in wing centering within the camera frames. Although the

roll-induced effect on span is minor, it becomes measurable for large off-centering angles, as expected from the in-plane shear distortion introduced by roll misalignment.

Despite the significant improvement after optimization, the mean epipolar error remains non-negligible. In some cases, vertical disparity decreased by more than two orders of magnitude, yet did not reach ideal stereo calibration levels. This residual error is consistent with the assumption that only the rotation matrix \mathbf{R} was corrected, while the translation \mathbf{t} remained fixed. Ignoring small translational effects—caused by elastic deformation of the Γ -bar—inevitably introduces a residual misalignment, especially for distant markers.

Additional uncertainty arises from marker visibility and subpixel detection. During flight, several markers were partially occluded or shadowed by the wing, reducing contrast and affecting centre detection accuracy. Although the subpixel localization algorithm (see Chapter 4) performs well for clearly visible markers, partial occlusion introduces random errors in the range of 1–2 px, which may explain the remaining epipolar deviations. Including translation optimization would further improve accuracy, but at the cost of tripling the search space dimensionality, thereby increasing computational complexity by approximately $\mathcal{O}(n^3)$.

Given the small angular deviations observed ($< 2^\circ$) and the already satisfactory results, the correction of rotational parameters alone was deemed sufficient for this study.

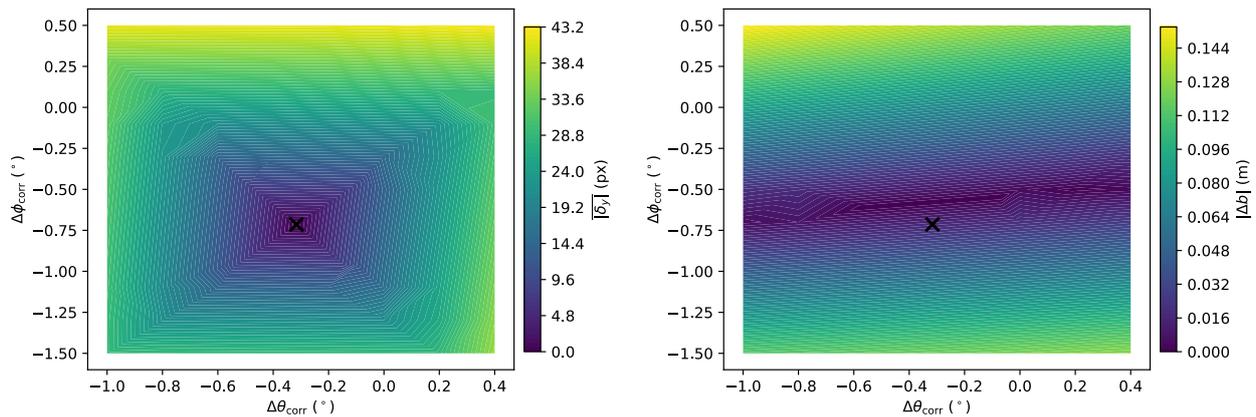


Figure 8.6: Parameter slice through the optimized yaw correction $\Delta\psi_{\text{corr}}^*$. (Left) Mean vertical disparity $|\delta_y|$ as function of roll and pitch corrections; (Right) corresponding span deviation Δb compared to the UWB reference. Roll and pitch corrections mainly affect vertical alignment, while span variations are dominated by yaw.

Automated stereoscopic calibration correction

While the manual optimization framework is effective for isolated frames, it becomes impractical for continuous video analysis at frame rates around 30 fps. Therefore, an attempt was made to automate the calibration correction using correlation analysis between the optimized angular corrections ($\Delta\psi_{\text{corr}}$, $\Delta\phi_{\text{corr}}$, $\Delta\theta_{\text{corr}}$) and relevant flight parameters.

The underlying assumption was that the bending and twisting of the Γ -bar are driven by the combined influence of bridle forces and aerodynamic loads acting on the kite. The deformation of the bar depends both on the magnitude of the forces transmitted through the bridles and on their orientation, which varies dynamically with steering input and depower setting. In addition, aerodynamic drag and lift contribute to the total load, scaling approximately with the square of the apparent wind speed (V_a^2). During turning manoeuvres, the centrifugal component of the aerodynamic load also increases, further influencing the effective bridle tension and thus the bending moment on the bar.

To explore these dependencies, a correlation analysis was performed between the optimized correction angles and various flight parameters, combining up to two predictors per test. The following variables

were considered representative of the physical loads acting on the bar: tether force F_T , depower setting, steering input u_s , apparent wind speed V_a (and V_a^2), turn rate, and the relative orientation of the front bridles. These parameters collectively describe both the aerodynamic environment and the mechanical boundary conditions experienced by the camera rig.

Despite these physically motivated hypotheses, the analysis did not reveal any clear or consistent dependencies. This suggests that the deformation of the Γ -bar arises from a complex, nonlinear coupling between aerodynamic forces, bridle geometry, and structural stiffness, which was not captured by simple regression model. Cross-interactions between yaw, roll, and pitch deformations likely further obscure any direct relationships.

Although a more advanced, data-driven predictive model could in principle be developed using machine learning or detailed structural simulations, this was beyond the scope of the current study. Instead, a pragmatic semi-automated approach was adopted. Manual optimizations were performed at regular time intervals (e.g. every second), after which the angular corrections were linearly interpolated between these points. Given the small angular deviations and gradual evolution of flight states, this interpolation was found to be sufficiently accurate when considering small time intervals. In practice, this approach requires approximately ten manual optimizations for a ten-second manoeuvre, after which all intermediate frames can be automatically corrected for subsequent photogrammetric reconstruction, using the photogrammetry algorithm detailed in Chapter 4.

Adjusted Accuracy

The application of angular corrections inevitably affects the effective accuracy of the photogrammetric reconstruction. While the baseline accuracy analysis in Chapter 4 applies to a rigid stereo rig, the present setup introduces additional uncertainty due to the flexible Γ -bar. To estimate the adjusted accuracy $\epsilon_{p,adj}$, the reconstructed lengths of the carbon-reinforced struts were used as an internal consistency check.

Since the struts are structurally rigid, their true lengths remain constant throughout flight. Therefore, deviations in the reconstructed strut lengths directly represent residual geometric error after correction. The distances between the outermost markers on each strut are known and summarized in Table 8.3, note that this does not correspond to actual strut length.

Table 8.3: Reference distances between outermost markers on the carbon-reinforced struts used for adjusted accuracy estimation.

Strut (from centre)	True Distance (m)
1	2.02
2	1.92
3	1.53
4	1.14

For each strut where both outer markers were detected, the reconstructed 3D distance was compared to the true value, and the absolute deviation was used to estimate $\epsilon_{p,adj}$. This metric reflects the combined residual error arising from imperfect rotational correction, unmodelled translation, and remaining image detection noise. Across all analysed tests, $\epsilon_{p,adj}$ was found to be of the same order of magnitude as the accuracy determined under static calibration conditions (see Chapter 4), indicating that the post-processing correction effectively restores the geometric integrity of the reconstruction despite the small mechanical deformations of the camera rig. The adjusted accuracy analysis does not allow for distinguishing in-plane and out-of-plane errors.

8.2.2 (ii) Transformation to a wing-fixed reference frame

After the stereo geometry has been corrected for bar bending, all triangulated marker coordinates are still expressed in the camera frame of the left camera. The camera frame has its z -axis normal to the image plane, x horizontal and y vertical. To prevent rigid-body motion of the camera rig (or the entire system) from being misinterpreted as kite deformation, all reconstructed points are transformed to a wing-fixed coordinate frame before further analysis. This also enables a direct comparison of wing shapes for different manoeuvres.

The wing-fixed frame is constructed from the set of markers on the two centre struts, here denoted as strut 3 (centre-left) and strut 4 (centre-right).¹ For each frame, let $\{\mathbf{p}_{3,c}^{(j)}\}_{j=1}^{N_3}$ and $\{\mathbf{p}_{4,c}^{(j)}\}_{j=1}^{N_4}$ be the 3D camera-frame coordinates of all markers on strut 3 and strut 4, respectively.

The origin of the wing-fixed frame is placed midway between the centroids of the two centre struts,

$$\bar{\mathbf{p}}_{3,c} := \frac{1}{N_3} \sum_{j=1}^{N_3} \mathbf{p}_{3,c}^{(j)}, \quad \bar{\mathbf{p}}_{4,c} := \frac{1}{N_4} \sum_{j=1}^{N_4} \mathbf{p}_{4,c}^{(j)}$$

$$\mathbf{o}_w := \frac{1}{2} (\bar{\mathbf{p}}_{3,c} + \bar{\mathbf{p}}_{4,c}).$$

A right-handed orthonormal basis $\{\hat{\mathbf{x}}_w, \hat{\mathbf{y}}_w, \hat{\mathbf{z}}_w\}$ is then constructed as follows. First, unit direction vectors $\hat{\mathbf{d}}_3$ and $\hat{\mathbf{d}}_4$ are estimated along strut 3 and strut 4 (from a simple least-squares line fit to their markers). The chordwise axis is taken as the normalised average of these two strut directions,

$$\hat{\mathbf{x}}_w := \frac{\hat{\mathbf{d}}_3 + \hat{\mathbf{d}}_4}{\|\hat{\mathbf{d}}_3 + \hat{\mathbf{d}}_4\|}.$$

A raw spanwise direction is defined by the line between the two strut centroids,

$$\hat{\mathbf{y}}_{\text{raw}} := \frac{\bar{\mathbf{p}}_{3,c} - \bar{\mathbf{p}}_{4,c}}{\|\bar{\mathbf{p}}_{3,c} - \bar{\mathbf{p}}_{4,c}\|}.$$

The normal vector is then the cross product of the chordwise and raw spanwise axes,

$$\hat{\mathbf{z}}_w := \frac{\hat{\mathbf{x}}_w \times \hat{\mathbf{y}}_{\text{raw}}}{\|\hat{\mathbf{x}}_w \times \hat{\mathbf{y}}_{\text{raw}}\|},$$

and the spanwise axis is orthogonalised as

$$\hat{\mathbf{y}}_w := \frac{\hat{\mathbf{z}}_w \times \hat{\mathbf{x}}_w}{\|\hat{\mathbf{z}}_w \times \hat{\mathbf{x}}_w\|}.$$

Collecting these unit vectors as columns yields the rotation matrix from camera to wing axes,

$$\mathbf{R}_{wc} := [\hat{\mathbf{x}}_w \quad \hat{\mathbf{y}}_w \quad \hat{\mathbf{z}}_w] \in \mathbb{R}^{3 \times 3}.$$

Any camera-frame point \mathbf{p}_c is then mapped to the wing frame via the rigid transform

$$\mathbf{p}_w = \mathbf{R}_{wc}^\top (\mathbf{p}_c - \mathbf{o}_w).$$

If some centre-strut markers are momentarily not reconstructed in a given frame, their positions are first interpolated on the common time axis (see Section 8.1.2) before constructing \mathbf{o}_w and \mathbf{R}_{wc} , which avoids spurious jumps in the transformed wing shapes. Figure 8.7 shows a graphical depiction of both the camera and the wing fixed reference frame.

¹Strut numbering follows from counting struts from left to right, starting at 0. This results in strut numbered 3 and 4 being the centre struts.

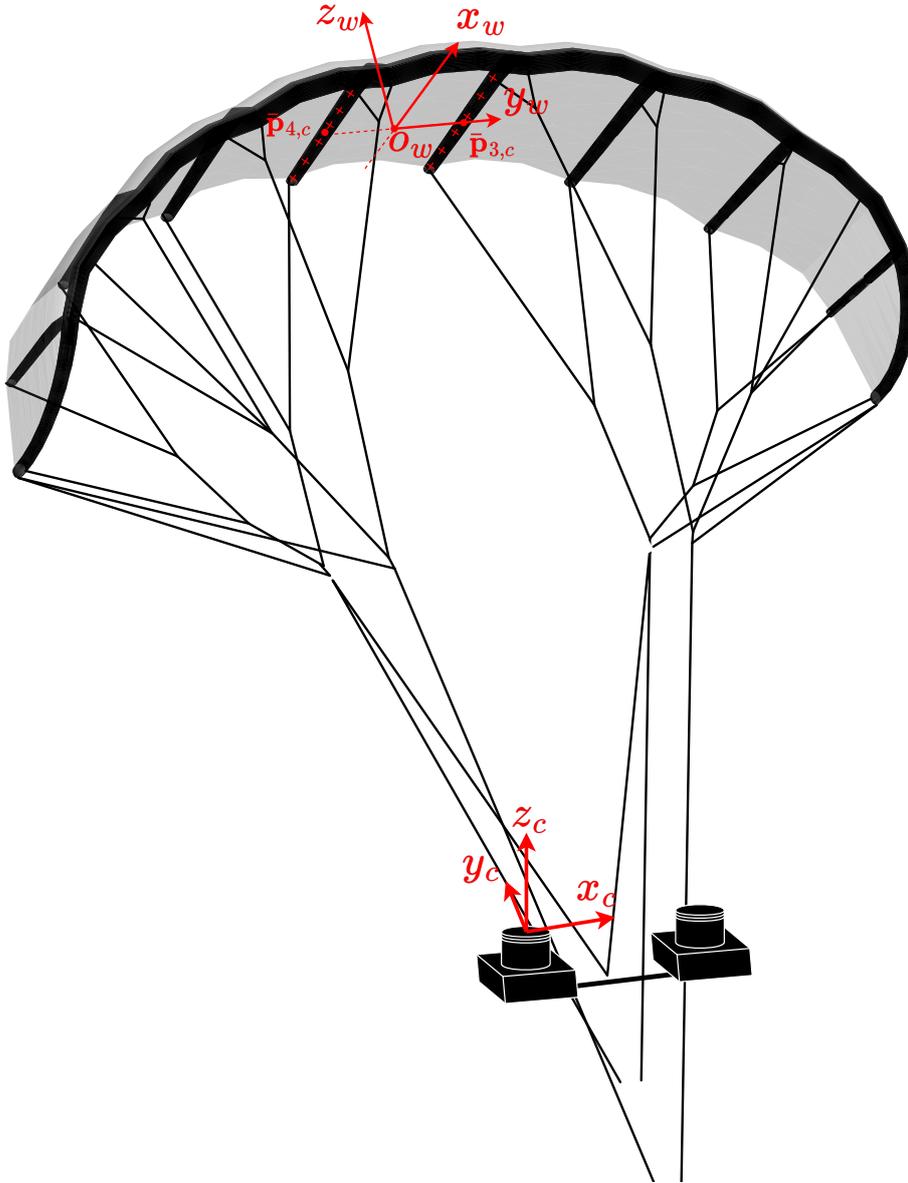


Figure 8.7: Camera and wing reference frames. The dense sets of markers on the two centre struts (strut 3 and strut 4) define a wing-fixed frame at each instant: the origin \mathbf{o}_w is located midway between the strut centroids, the x -axis $\hat{\mathbf{x}}_w$ follows the average direction of the two centre struts, the y -axis $\hat{\mathbf{y}}_w$ is spanwise, and the z -axis $\hat{\mathbf{z}}_w$ is normal to the wing plane.

8.2.3 (iii) Outlier rejection in the 3D marker cloud

Despite careful stereo matching and filtering, some spurious points remain in the triangulated marker cloud, typically originating from mismatched correspondences or partial occlusions. These outliers are removed using a simple radius-based nearest-neighbour criterion applied per frame in 3D space.

Let $\{\mathbf{p}_i\}$ be the set of reconstructed points in the wing frame for a given frame. A point is retained only if it has at least one neighbour within a configurable distance r_{lim} :

$$\text{keep } \mathbf{p}_i \iff \min_{j \neq i} \|\mathbf{p}_i - \mathbf{p}_j\|_2 \leq r_{\text{lim}}.$$

Points failing this test are discarded as outliers. The threshold r_{lim} is tuned per flight and chosen smaller than the typical spacing between the LE and strut markers, so that isolated, incorrectly

triangulated detections are removed while coherent clusters on the LE and struts are preserved (REQ-DPV-07).

8.2.4 (iv) Wing shape reconstruction

After transformation to the wing-fixed frame and outlier rejection, the remaining markers form a sparse but coherent sampling of the inflatable wing structure. Two classes of markers are distinguished: markers on the LE and markers on the struts. Using the LE/strut classification obtained from the photogrammetry algorithm discussed in Section 4.6, the wing shape is reconstructed by fitting geometric models to the corresponding 3D points.

Leading-edge reconstruction

Let $\{\mathbf{p}_k\}_{k=1}^{n_{LE}} \subset \mathbb{R}^3$ be the LE points expressed in the wing frame. The LE is represented by a polyline that connects all these points in an order such that the total length of the polyline is minimized. This follows the idea that the physical LE is, by design, the shortest curve that connects all LE markers in order, i.e. a shortest Hamiltonian path without returning to the start [137]:

$$\min_{\pi \in S_{n,LE}} L(\pi) = \sum_{k=1}^{n_{LE}-1} \|\mathbf{p}_{\pi(k+1)} - \mathbf{p}_{\pi(k)}\|_2,$$

where π is a permutation of the indices and $S_{n,LE}$ the set of all permutations.

In practice this combinatorial problem is solved approximately. An initial polyline is constructed by iteratively connecting each marker to the closest not-yet-visited marker (nearest-neighbour initialisation) and then refined using 2-opt edge swaps to shorten the polyline. To reduce computational cost, candidate start–end pairs are restricted to the subset of LE markers with the lowest camera-frame z -coordinate (here the eight lowest points), and the path with the smallest total length among these trials is retained. The resulting polyline consists of straight segments between successive LE markers and bends only at the markers themselves. Given the marker density used in the experiments, this provides an accurate approximation of the true LE shape.

Strut reconstruction

Let $ce = (ce_x, ce_y)$ denote the 2D image coordinates of markers centroid classified as strut markers in a given frame. These strut markers are first partitioned into eight vertical bands in image space using a horizontal spacing s . Sorting the candidates ce by pixel x -coordinate, clusters C_k are grown left-to-right by comparing each point to the first point of the current cluster:

$$\text{assign } ce \text{ to cluster } C_k \iff |x_i - x_{\text{first}(C_k)}| < s \text{ (px)},$$

with spacing threshold s (here $s = 130$ px), resulting in eight clusters C_k (with $k = 0, \dots, 7$). The result is shown in Fig. 8.8.

Each strut is modelled as a straight line, justified by the internal carbon-fibre tubes that maintain rigidity during flight. Let $\{\mathbf{p}_j\}_{j=1}^m \subset \mathbb{R}^3$ be the wing-frame points associated with one strut. The centroid is

$$\bar{\mathbf{p}} = \frac{1}{m} \sum_{j=1}^m \mathbf{p}_j,$$

and the centred point matrix is

$$\mathbf{Q} = \begin{bmatrix} \mathbf{p}_1 - \bar{\mathbf{p}} \\ \mathbf{p}_2 - \bar{\mathbf{p}} \\ \vdots \\ \mathbf{p}_m - \bar{\mathbf{p}} \end{bmatrix}.$$

A singular value decomposition (SVD) of \mathbf{Q} ,

$$\mathbf{Q} = \mathbf{U}\mathbf{S}\mathbf{V}^\top,$$

yields the right singular vectors in \mathbf{V} . The first right singular vector (the first column of \mathbf{V}) defines the unit direction vector $\hat{\mathbf{v}}$ of the best-fit line. The strut is then represented parametrically as

$$\mathcal{L}(t) = \bar{\mathbf{p}} + t\hat{\mathbf{v}}, \quad t \in \mathbb{R}.$$

This line minimises the sum of squared orthogonal distances of all points to the line:

$$\min_{\hat{\mathbf{v}}, \bar{\mathbf{p}}} \sum_{j=1}^m \left\| (\mathbf{I} - \hat{\mathbf{v}}\hat{\mathbf{v}}^\top) (\mathbf{p}_j - \bar{\mathbf{p}}) \right\|_2^2.$$

This total-least-squares formulation is appropriate because triangulation noise affects all three spatial coordinates.



Figure 8.8: Marker indexing on the wing during flight. The distinction between the leading-edge (LE) markers and the strut markers is indicated, with the strut markers numbered from 0 to 7. Some markers remain unidentified, either because they do not pass the detection filters or because they are occluded in either of the camera views.

Taken together, the LE polyline and the fitted strut lines form a sparse wireframe representation of the inflatable wing. Applied independently to each frame, this procedure yields a time-resolved 3D reconstruction of the wing deformation throughout flight.

This chapter discusses the results obtained in the flight test, with a particular focus on how the observed behaviour relates to previously stated claims in the literature regarding depower and steering.

To address this objective, a set of representative static flight cases is analysed. These cases consist of pairs of closely spaced time instants (typically separated by only a few seconds) selected immediately before and during a change in control input. The moments are not time-averaged, but treated as instantaneous snapshots of the system state. Selection was restricted to time instants for which complementary UWB data were available, such that the required geometric corrections could be applied consistently. For depowering, the comparison is made between a fully powered state and depowered state, in compliance with settings for the V3 kite, occurring within the same transition. For steering, cases are chosen where the steering input changes rapidly, producing a clear difference in turn rate over a short time interval. In all cases, the selected instants correspond to pronounced changes in the measured variables, ensuring that the resulting differences can be attributed directly to the applied control action.

The chapter starts with comparing the geometry of the wing in flight to the CAD geometry used in aerodynamic analyses in Section 9.1. Afterwards it evaluates two behaviours, billowing and twist, for different flight phases (Section 9.2). Lastly, it quantitatively evaluates what occurs when depowering and steering during flight, in Section 9.3 and Section 9.4 respectively.

9.1 In-flight vs. CAD geometry

A comparison between the undeformed CAD geometry and the reconstructed powered in-flight geometry provides insight into how the flexible LEI structure deforms under aerodynamic loading. The CAD model used here corresponds to the TU Delft V3A kite reference shape, which has been used extensively in earlier aerodynamic studies. Leuthold [39] relied on this rigid geometry to develop a quasi-steady multiple-wake vortex lattice method, while Demkowicz [23] and Lebesque [25] applied the same model in steady-state RANS simulations to analyse three-dimensional flow separation, spanwise transport, and the influence of struts and sideslip. Poland et al. [138, 139] used a CAD-based 1:6.5 rigid scale model of the V3 kite in a wind tunnel to analyse flow behaviour and wing loading. The CAD representation therefore forms the baseline against which the true in-flight shape can be evaluated.

The comparison in Fig. 9.1 shows that in the wing reference frame, the most prominent difference between the CAD and in-flight geometry is the effective span as a result of billowing. The powered in-flight shape exhibits a clear reduction in span, and the outer struts rotate progressively with respect to the centre. This behaviour indicates substantial billowing of the canopy between struts, a deformation mode that is low and fixed in the undeformed CAD configuration.

Billowing pulls the trailing edges of the struts closer together while the spacing between the leading edges remains largely unchanged due to the pressure of the inflated LE tube. This contraction of the trailing edge causes a rearward displacement of the LE tube near the tips, which alters the local inflow angle and modifies the spanwise loading distribution. The bottom and front views also show a pronounced inward displacement of the trailing-edge tips. Although mainly driven by billowing, this effect also reflects a redistribution of tension within the inflated body.

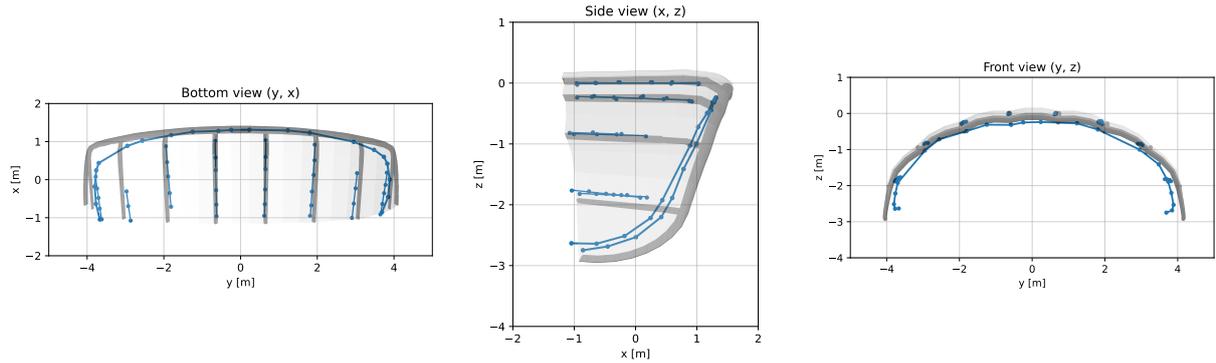


Figure 9.1: Comparison between the powered in-flight kite geometry (blue) and the undeformed CAD representation (grey).

The side-view comparison reveals a symmetric downstream deflection of the wingtips, meaning the in-flight tips bend backwards under load. This is consistent with drag-induced deformation of the system, resulting in a symmetric change in pitch distribution along the span. This affects the local angle of attack α and is therefore relevant for both aerodynamic efficiency and dynamic stability.

Contrary to assumptions made in the aeroelastic modelling work of Cayon [140], the powered wing does not flatten under load. Instead, both the leading-edge and trailing-edge spans contract, and the planform becomes more curved. The combined effects of billowing and drag-induced pitch alteration pulling the tips backwards lead to an in-flight geometry that is more curved and differs systematically from the undeformed CAD representation.

A quantitative comparison of the span is provided in Table 9.1. The effective span decreases by 10.73% in the powered in-flight condition relative to the CAD model. This constitutes a substantial deviation for aerodynamic analysis and is primarily driven by aeroelastic deformation and bridle load redistribution. The stronger curvature and increased effective anhedral observed in flight are consistent with the contraction induced by billowing.

This contraction also reduces the wing’s projected area. Using the bottom-view projection of the reconstructed geometry, the projected area decreases by 12.15% compared to the CAD reference. To account for the fact that the outer markers are positioned approximately 0.1 m inboard of the true wing boundary, the polygon representing the projected outline was dilated by 0.1 m before computing its area. This corrected projected area better captures the true planform of the curved, in-flight wing and is therefore more representative for aerodynamic analysis. As this is an estimation, no uncertainty is defined. The decrease in projected area highlights that aerodynamic predictions based solely on rigid CAD data overestimate real flight conditions.

Table 9.1: Measured reduction in span between the undeformed CAD geometry and the powered in-flight geometry.

Case	Span b (m)	Projected area A (m ²)
CAD model	8.20	19.75
Powered	7.32	17.35
Difference	$-10.73 \pm 0.30\%$	-12.15%

9.2 Billowing and twist

This section discusses two deformation mechanisms that can be directly quantified from the photogrammetry results: billowing and twist. Both effects were introduced qualitatively in the previous section, but their quantitative evaluation is essential for characterising the in-flight geometry across different manoeuvres.

9.2.1 Billowing

Billowing is a global deformation mode in which the canopy deflects between the struts under aerodynamic loading. This behaviour, also referred to as the adaptive camber of LEI kites [39], is (nearly) absent and constant in the undeformed CAD configuration and becomes more pronounced as aerodynamic loading increases. Billowing can be quantified by evaluating the spanwise distances between trailing-edge markers on consecutive struts, as defined by the white segments in Fig. 9.2. A reduction in segment length indicates an increase in billowing, and vice versa.

The measured segment lengths are shown in Fig. 9.3. Because the markers are positioned slightly inboard of the true trailing edge, the actual strut-to-strut distances are marginally smaller than the measured values. For the CAD model, this offset is negligible, as the struts are parallel to the streamwise direction. Given the small magnitude of this effect, it is considered negligible for the present analysis.

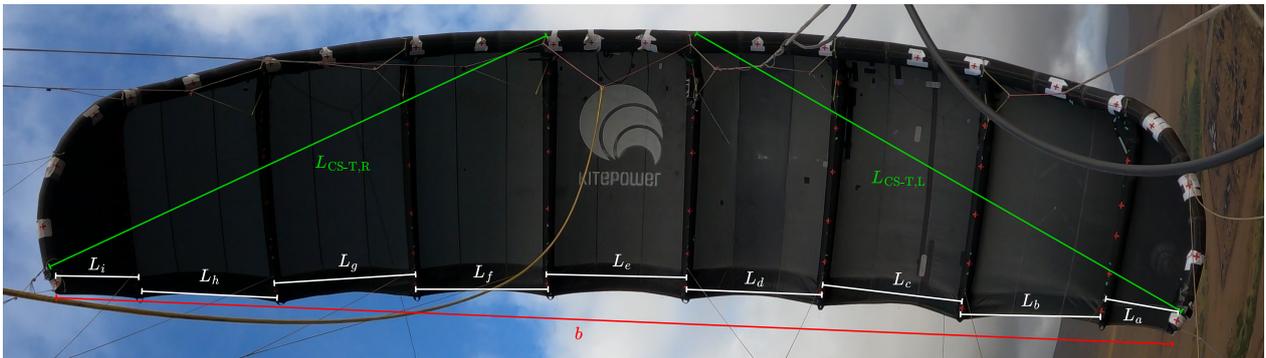


Figure 9.2: Segments considered for billowing analysis (in white). This video still is taken during a right turn.

Across nearly all segments, the in-flight distances are shorter than those of the CAD model, confirming that the canopy contracts under load and forms a spanwise arc at the trailing edge. Only the outer segments—between the last struts and the tips—show smaller or inconsistent reductions. These deviations are likely linked to both lower aerodynamic loading near the wing tips and reconstruction uncertainty; the average positional error $\epsilon_{p,adj,avg}$ is of the same order as the small differences observed in this region.

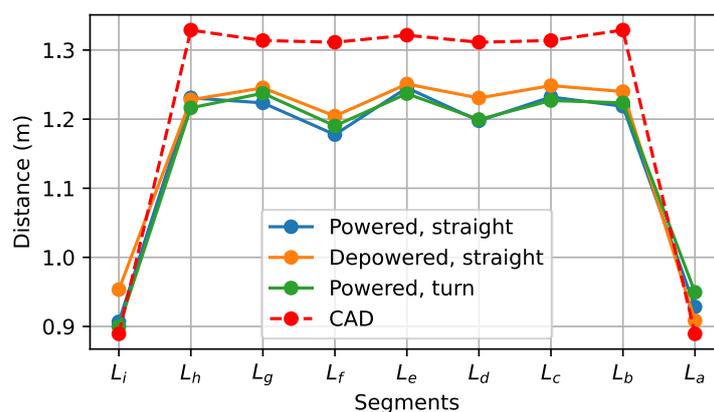


Figure 9.3: Length of kite segments, indicating billowing for all flight cases.

Although all in-flight cases show reduced strut-to-strut distances compared to the CAD geometry, the magnitude of the reduction varies only modestly between manoeuvres and is not perfectly symmetric across the span. These variations fall within the typical reconstruction uncertainty and should therefore not be overinterpreted. A more robust measure is the average reduction in segment length across the span, as averaging significantly reduces noise. The resulting mean percentage change

relative to the CAD configuration, $\Delta\overline{L}_n$, is shown in Table 9.2. Due to the complicated nature of the billowing segments, which are inversely correlated, the uncertainty calculation required a complex model; too complex in fact for the time at hand.

Table 9.2: Difference in strut-to-strut distance with respect to the CAD model for several flight states.

Flight case	$\Delta\overline{L}_n$ wrt CAD (%)	F_T/g (kg)
Powered, straight	-5.23	325.08
Powered, turn	-5.00	283.73
Depowered, straight	-3.86	46.94

These results clearly show that billowing increases with aerodynamic loading. Even in the depowered condition, where the tether force is relatively small, the canopy still contracts compared to the CAD geometry because the membrane carries aerodynamic load once the kite is airborne. Since the tether force grows with aerodynamic force F_a (see Fig. 2.3), larger flight loads correspond to stronger billowing. However, the relationship is not linear: it depends on flight speed, canopy design, initial membrane tension, and the unloaded geometry.

Finally, it should be noted that during the experiment the KCU was relatively heavy and the ambient wind speed was low, resulting in a slack tether and comparatively small aerodynamic forces. Despite this, billowing remained clearly measurable, illustrating that it is an intrinsic aeroelastic behaviour of LEI kites rather than a secondary effect.

9.2.2 Twist angle

The twist angle τ quantifies how each strut rotates relative to the local orientation of the LE in the zx -plane of the camera coordinate frame. By definition, the centre of the wing has $\tau = 0^\circ$, such that twist is always measured relative to the central region between the two middle struts. This avoids contamination by global roll or pitch motions of the entire kite. The definition is conceptually similar to the local geometric angle of attack discussed by Cayon [140], but with an important distinction: because all values are referenced to the centre where $\tau = 0^\circ$, the quantity represents shape change rather than the absolute geometric angle of attack.

To illustrate, if the entire kite pitches up by n° , the centre region experiences a n° increase in geometric angle of attack. However, the twist distribution remains unchanged, because each strut is still oriented identically relative to the central reference. Determining a global geometric angle of attack is challenging for LEI kites due to the difficulty of defining a unique reference chord or “front” in the deformed in-flight configuration. The twist angle therefore provides a more robust metric for describing spanwise deformation. A positive τ indicates that the TE rotates inward toward the KCU, increasing the local pitch, whereas a negative τ corresponds to the TE rotating outward. This definition is visualised in Fig. 9.4(a).

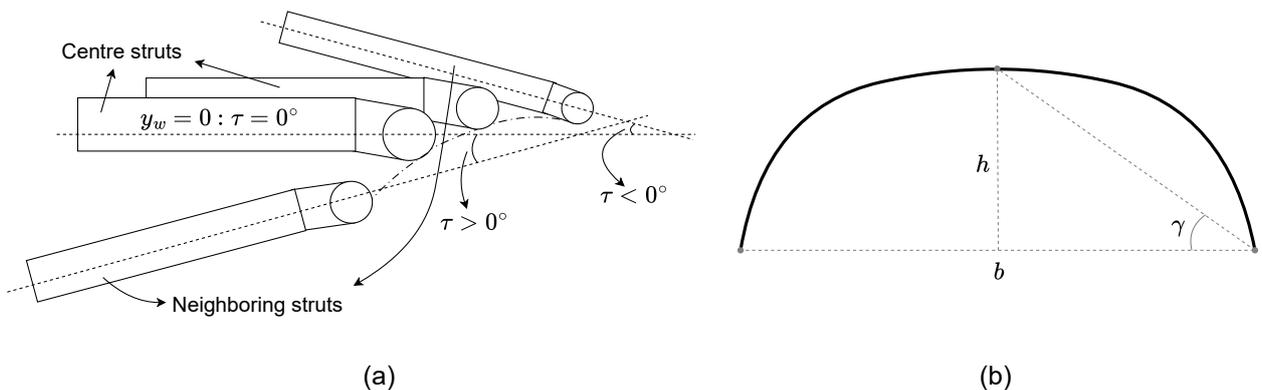


Figure 9.4: Definition of (a) twist angle τ and (b) anhedral angle γ .

The CAD geometry exhibits almost no twist, with only slight negative values near the tips, indicating a small outward rotation of the TE in the unloaded configuration. In contrast, the powered and depowered in-flight configurations both show a symmetric increase in twist magnitude toward the tips. Minor asymmetries between the left and right sides are expected due to measurement uncertainty, imperfections in node-based orientation extraction, and the inherently unsteady nature of real flight. Even in “straight” flight, figure-of-eight manoeuvres rarely produce a perfectly symmetric condition.

The predominantly positive twist angles confirm that the TE is displaced inward, consistent with the billowing-induced contraction of the trailing-edge region discussed earlier. Small negative values appear on the right side of the wing in the depowered-state measurements, likely caused by slight asymmetric flight or an apparent-wind angle offset. These variations fall within the expected experimental scatter. Nevertheless, the overall trend is clear and robust: the in-flight wing exhibits more twist than the undeformed CAD configuration, especially toward the tips. This observation aligns with the aeroelastic deformation predicted in the model of Cayon [140].

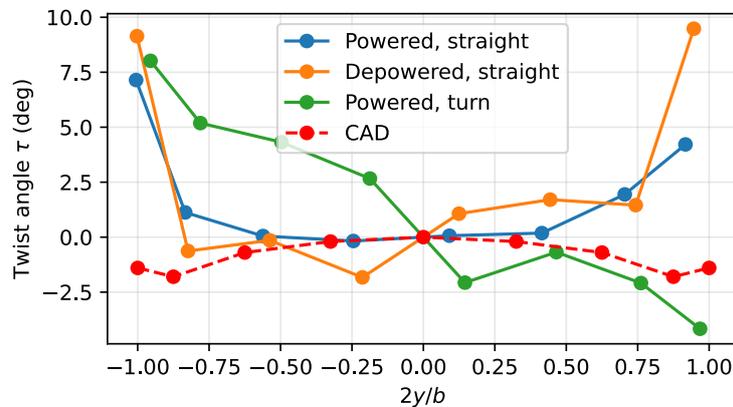


Figure 9.5: Spanwise twist angle τ for several flight states.

The largest variation in twist occurs during turning flight. In Fig. 9.5, the convention $2y/b > 0$ corresponds to the left side of the wing, and $2y/b < 0$ to the right. The data shown correspond to a right turn, meaning the pulled side of the wing appears at $2y/b < 0$ and the released side at $2y/b > 0$. The twist distribution shows a clear twist reduction on the released side and an increase on the pulled side. This differential twist is consistent with the steering mechanism described by Breukels [17]: increasing local angle of attack on the pulled side and decreasing it on the released side modifies the spanwise lift and drag distribution. The differential in drag between the tips produces a yawing moment that drives the turn. Nevertheless, for highly curved LEI wings, the effective local angle of attack at the tips is governed by both the twist angle and the geometric sideslip induced during steering. The true aerodynamic angle of attack experienced by each strut in a turn is therefore a combination of these effects, which moderates the extreme values suggested by twist alone.

9.3 (De)power deformation

Depowering is essential for reducing aerodynamic loading during retraction. The literature describes several structural and aerodynamic changes associated with this manoeuvre, as summarised in Chapter 2. The most commonly cited claims are:

- (i) *Depowering increases the curvature of the planform, i.e. the wing exhibits a larger anhedral angle γ [8, 34].*
- (ii) *Depowering reduces the aerodynamic angle of attack α_w [8, 16].*
- (iii) *Depowering reduces the span [34, 140].*

- (iv) *The combined effect of reduced span, increased curvature, and lower angle of attack reduces the projected area and thus lowers the tether force F_T during retraction [8, 34].*

To evaluate these claims, the powered and depowered geometries from the flight test were analysed. Both reconstructed shapes were transformed into the wing reference frame to allow a consistent geometric comparison. Global deformation is compared in the camera reference frame.

Planform comparison

The reconstructed powered and depowered shapes are shown in Fig. 9.6, with corresponding flight parameters listed in Table 9.3. The steering input was negligible for this segment of the trajectory, meaning that differences between the two shapes can be attributed primarily to depowering. The bottom view shows only a minor difference in the curvature of the leading edge. Because the LEI tube is stiff, substantial changes in LE curvature are not expected. At the trailing edge, however, the strut spacing increases slightly when depowering, consistent with the observed reduction in billowing, flattening the wing. At the same time, the depower tape lengthens, relaxing the bridle system, which tends to increase curvature. These competing mechanisms result in only a modest increase in overall planform curvature.

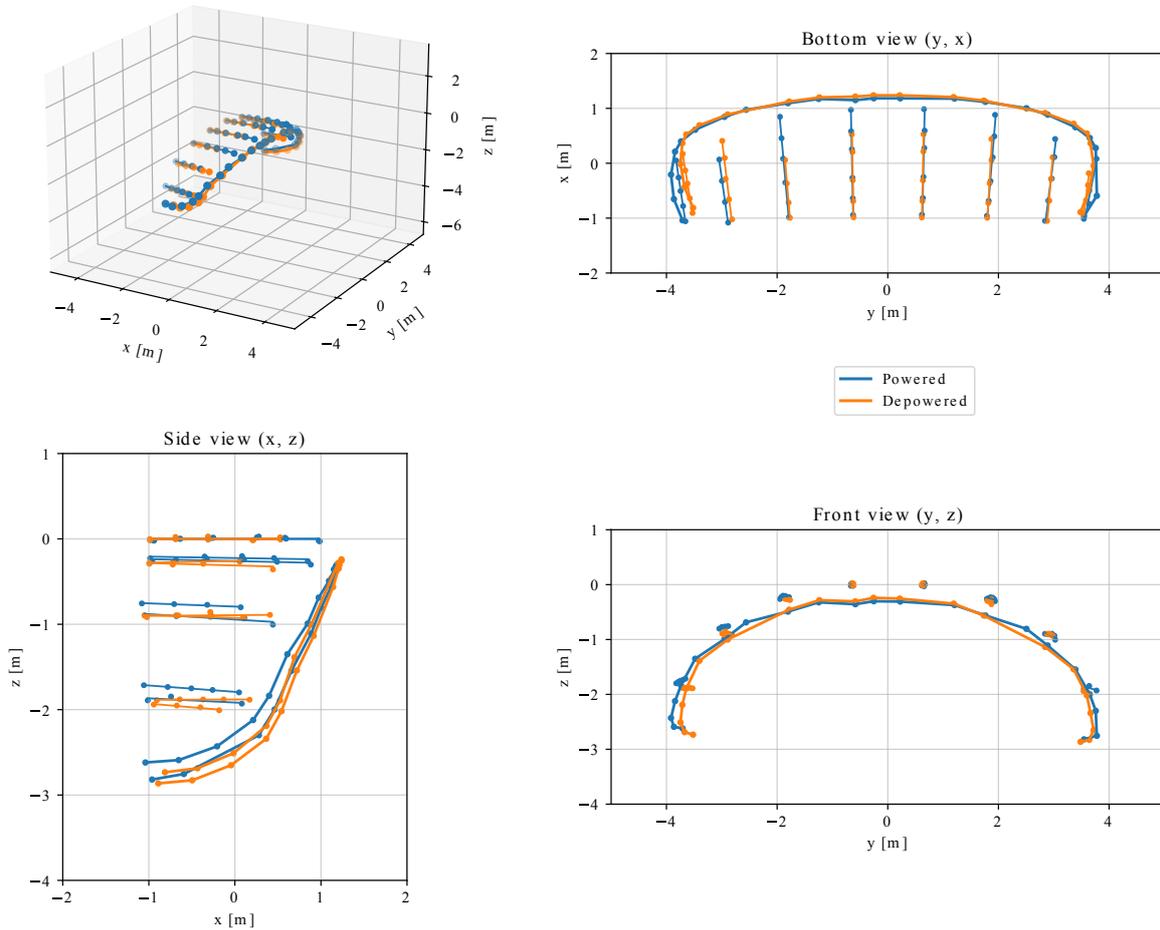


Figure 9.6: Characteristic geometry change between a powered and a depowered kite.

Table 9.3: In- and outputs corresponding to the two flight states provided in Fig. 9.6

State	u_{dp} (-)	u_s (-)	b_{UWB} (m)	V_a (ms^{-1})	α_{fl} ($^\circ$)	β_s ($^\circ$)	F_T/g (kg)	$\epsilon_{p,adj,avg}$ (cm)	$\epsilon_{p,adj,max}$ (cm)
Powered	0.4151	0.0028	7.32	16.75	10.9	-9	325.08	5.95	10.23
Depowered	0.5012	0.0014	7.12	10.56	19.4	-8	44.32	4.14	7.33

Although small, the measured change is in line with claim (i), meaning that the V3 kite does exhibit a slightly more curved planform when depowered, but the effect is weaker than implied in prior

modelling work. In earlier analytical models (e.g. the multi-plate model), billowing was typically underestimated [34, 140], resulting in overprediction of span reduction and increased curvature; the present data show a more moderate response.

Anhedral angle

The anhedral angle γ is illustrated in Fig. 9.4 and is defined as the angle between the centre of the leading edge and the TE tips in the zy -plane of the wing reference frame. The angle is directly related to the curvature of the wing. It is computed from:

$$\gamma = \arctan\left(\frac{h}{\frac{1}{2}b}\right),$$

where h is the vertical distance between the midpoint of the central leading-edge region and the midpoint of the tips, assuming symmetric deformation. Using this definition, the depowered state yields $\gamma_{\text{depowered}} = 38.17^\circ$, and the powered state $\gamma_{\text{powered}} = 36.62^\circ$, corresponding to an increase of 1.55° . Based on the average out-of-plane accuracy (for h) determined in Section 4.5, together with the UWB accuracy (see Section 5.3, for b), the uncertainty of γ is 0.67° (REQ-EXP-01c). This quantitatively confirms claim (i) to be true.

Wing pitch

To evaluate claim (ii)—that depowering reduces the angle of attack—the wing reference frame is not suitable, since it removes the global orientation of the kite. Instead, the pitch must be analysed in the camera frame. The side view in Fig. 9.7 shows that the depowered shape rotates forward relative to the powered shape, while the horizontal (y -) position remains similar. This indicates a true pitch change of the wing rather than a camera-induced artefact following from slack bridle lines. Another confirmation that the lines remain tensioned when depowering, indicating this is a true change in wing orientation, is the corrections that are still required during this flight phase. Nevertheless, the required bending and twist corrections are smaller than for the powered state, in line with the theory that the Γ -bar deformation occurs due to forces in the bridles.

The pitch angle was determined by averaging the inclination of the eight struts relative to the camera-frame y -axis. The pitch estimate is based on regression through many markers on eight

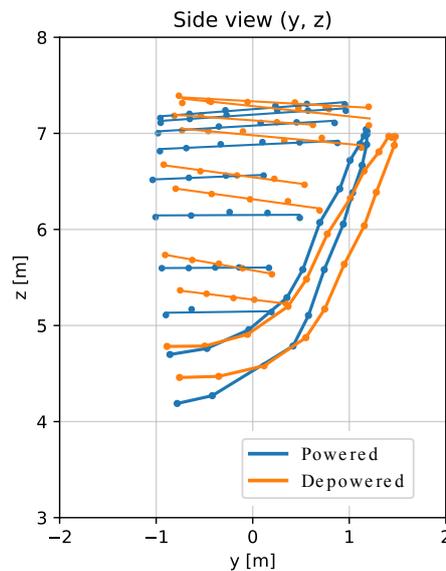


Figure 9.7: Side view of powered/depowered state in camera coordinate frame. The depowered state (orange) clearly exhibits a pitch forward with respect to the powered (blue) state.

different struts, reducing the effective uncertainty. The resulting uncertainty, based on the average out-of-plane accuracy determined in Section 4.5 was approximated to be 1.96° , more than REQ-EXP-01c. Nevertheless, this uncertainty is expected to be significantly lower, as the present approach neglects the redundancy inherent in the strut definition. Each strut is determined through a linear regression of multiple marker points, which effectively reduces the uncertainty of the estimated strut orientation and length. Accounting for this redundancy explicitly would require incorporating the varying number of identified markers per strut in each frame, which would considerably increase the complexity of the uncertainty analysis.

The measured pitch difference between powered and depowered states is $\Delta\theta_w = -8.69^\circ$, confidently confirming that depowering reduces pitch and therefore reduces angle of attack for the same inflow angle, validating claim (ii). The resulting $\Delta\theta_w$ is only 0.09° more than the input required for the EKF of Cayon et al. [141], further demonstrating the validity of this result.

Span reduction

The measured span values in Table 9.3 show that the depowered configuration has a span 2.73% smaller than the powered case. This is slightly lower than the 3–5% reductions previously estimated by Cayon and Poland [140, 34], who used both deformation models and single-camera photogrammetry to determine those values. In Poland’s single-camera photogrammetry approach, distortions in the video footage are corrected during post-processing. Nevertheless, because the reconstruction relies on a single viewpoint, the resulting measurements still show a relatively large uncertainty [34].

The magnitude of span reduction depends jointly on twist, billowing, and anhedral; in this dataset the change in twist is modest, and the dominant mechanism appears to be the increased anhedral, resulting in a small span reduction, since the reduction in billowing has an opposite effect. Despite this, the reduction agrees with claim (iii).

Projected area and tether force

Depowering reduces the projected area, though only slightly, but the reduction in angle of attack is substantial. This is reflected in the large decrease in mass-equivalent tether force—from 325 kg in the powered state to 44 kg in the depowered state—despite the modest changes in wing shape. For identical inflow conditions, a reduction in pitch directly reduces the aerodynamic angle of attack and thus the lift force acting on the kite. The data therefore support claim (iv). However, the drop in aerodynamic loading during depower is dominated by the reduction in angle of attack rather than by large structural changes in the planform. Nevertheless, the measured apparent wind speed also decreases when depowering, which further weakens the aerodynamic loading and therefore leads to an additional reduction in tether force (F_T). This reduction in (V_a) is consistent with crosswind-flight theory: in ideal crosswind operation, the kite’s lateral motion increases the apparent wind and thus the aerodynamic forces, typically yielding higher (V_a) and larger tether forces than in depowered, upward flight [7].

Summary of relevant findings

In summary, the kite geometry changes only moderately when depowered: anhedral increases slightly, billowing is reduced due to lower aerodynamic loading, and the span contracts by a few percent. The depowered shape therefore resembles the CAD geometry marginally more closely than the highly loaded powered shape. The dominant aerodynamic consequence of depowering is the reduction in pitch and thus angle of attack, which explains the large drop in tether force. All four claims from the literature are supported by the measurements, though the magnitude of geometric change—especially curvature and span—is smaller than in earlier, simplified models.

The key quantitative differences between powered and depowered flight are summarised in Table 9.4.

Table 9.4: Summary of relative findings when going from fully powered to fully depowered state. Uncertainties are based on the average photogrammetry and UWB accuracies, found in Chapter 4 and Chapter 5 respectively

Changing variable	Symbol	Value	Uncertainty
Depower constant	Δu_{dp}	+0.086	Unknown
Anhedral	$\Delta\gamma$	+1.5°	$\pm 0.67^\circ$
Wing pitch	$\Delta\theta_w$	-8.69°	$\pm 1.96^\circ$
Span	Δb	-2.73%	$\pm 0.60\%$
Strut TE distance	$\Delta \bar{L}_n$	+1.45 %	ND

9.4 Steering-induced deformations

Steering an LEI kite involves creating an asymmetric canopy deformation by adjusting the steering-line tensions. In this work, turning direction is defined in the wing-fixed reference frame from the perspective of an observer “riding” on the wing and aligned with the direction of flight (i.e., the vehicle-dynamics convention, not an external view from below). With this convention and the right-hand rule, a right turn corresponds to a negative rotation about the wing (z)-axis, and a left turn to a positive rotation about the wing (z)-axis. Defining left/right in this way removes ambiguity from camera viewpoint and avoids apparent left–right inversions in certain views (e.g., when the wing is viewed from behind, the left tip still appears on the left). In a right turn, the right tip is the pulled tip, and the left tip is the released tip. The results for a left turn are shortly touched upon in Appendix H. The literature identifies three primary deformation mechanisms associated with turning:

- (i) *Asymmetric deformation increases the local angle of attack α_{local} on the pulled tip and reduces it on the released tip, generating a yawing moment [17].*
- (ii) *Increased tension on the pulled steering line expands the effective aerodynamic area on that side, while decreasing it on the released side, resulting in a roll moment [17].*
- (iii) *Steering induces asymmetric out-of-plane deformation: the pulled tip bends forward and the released tip bends backward, further increasing asymmetry in the loading, thereby producing a yawing motion [17] [140].*

These mechanisms describe the aerodynamic and structural asymmetries that facilitate the turning of a LEI kite. To validate these claims, a quantification method is defined, which is illustrated in Fig. 9.8. In this case a right turn is illustrated, corresponding to negative yaw and positive roll moments. To quantify the motion of the tips, two deformation angles are defined relative to the geometry in straight flight:

- The **front-view shear angle** $|\Delta\xi|$, representing vertical motion in the zy_w -plane (roll-related deformation).
- The **bottom-view shear angle** $|\Delta\zeta|$, representing horizontal motion in the xy_w -plane (yaw-related deformation).

In Fig. 9.8(i), the drag differential due to the differential angle of attack at the tips is shown. This force difference results in a global negative yawing motion. Despite the lower drag, the released tip moves downstream relative to the pulled tip, as it is free to move due to the slack line. Claim (ii) is best evaluated from the front view, where a change in effective aerodynamic area and an offset in lift loads on both sides cause a positive roll moment. Claim (iii) works on the same principle but is based on the differential lift at the tips in the x_w, y_w plane. The released side is more rearward than the pulled side, which generates an additional negative yawing moment on the wing.

Observed planform deformation

The reconstructed geometries for both straight flight and a right turn are shown in Fig. 9.9. When both shapes are transformed into the wing reference frame by aligning the center struts, a clear

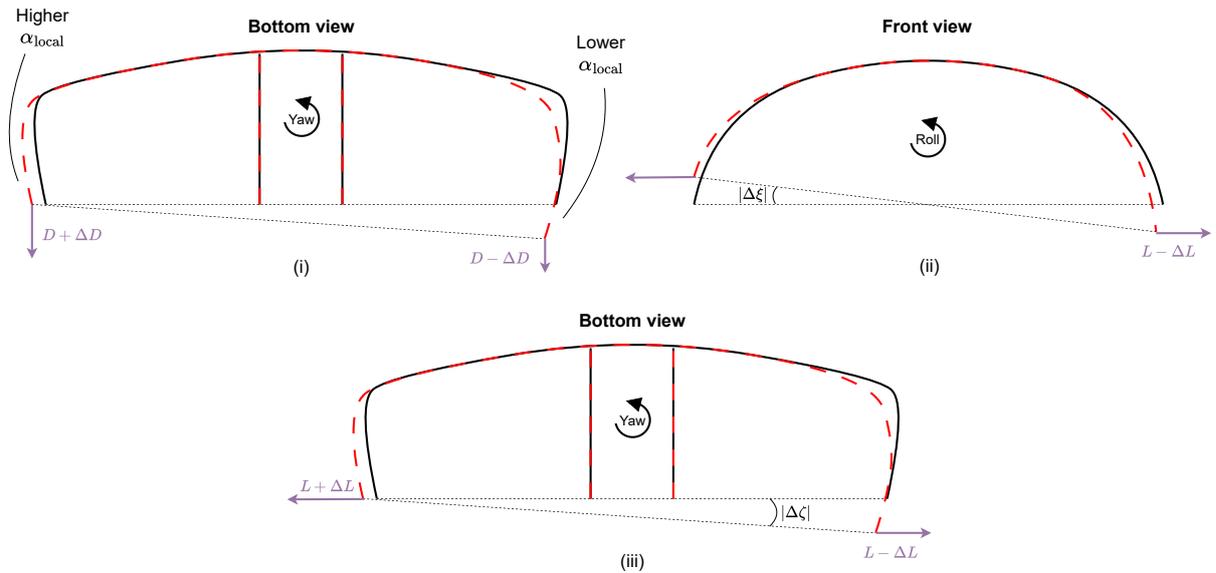


Figure 9.8: Shape change for the claims made, all in the wing-fixed reference frame. The identifiers correspond to the claims: (i) and (iii) result in a yaw motion (bottom view), and (ii) causes a roll motion (front view). The variable local angle in (i) is a consequence of the variable twist (see Fig. 9.5).

spanwise asymmetry becomes evident: the pulled tip moves upward and slightly upstream, while the released tip moves downward and rearward. This indicates that the asymmetric steering load induces not only a rotation of the wing but also a measurable redistribution of the canopy shape.

Flight parameters corresponding to straight flight and the right turn are presented in Table 9.5. The steering input u_s transitions from near zero in straight flight to a significant negative value during the turn, corresponding to the shortening of the right steering line and loosening of the left one. The tether force increases significantly in the turning segment, reflecting the higher aerodynamic loading. This change in tether force is not consistent in turns and is situation-dependent.

Local angle of attack and twist

As detailed in Section 9.2.2, the twist angle τ quantifies the rotation of each strut relative to the central leading-edge direction. Positive τ corresponds to inward rotation of the trailing edge toward the KCU, which increases the local pitch and the local geometric angle of attack. Because τ is measured relative to the wing centre, it isolates shape change rather than global pitch. During the turn, the twist distribution becomes asymmetric: the pulled tip experiences increased positive twist, while the released tip shows reduced or negative twist. This directly modifies the local angle of attack. Although sideslip β_s moderates the effective values at the tips, the spanwise asymmetry in τ dominates, as sideslip is constant over the span and lower than the extreme τ values.

Another turn-related effect is a spanwise variation in apparent wind speed. During a (yawing) turn, points along the span have different radial distances to the centre of rotation. For a given turn rate, the local tangential velocity therefore increases with radius, so the released (outer) side experiences a slightly higher apparent wind speed V_a than the pulled (inner) side. This introduces a spanwise differential in the apparent wind speed V_a . For the V3 kite, the wingspan is approximately 7–7.5 m, while the turn radius is estimated to be $r \approx 23$ m. Let the tangential velocity associated with the yawing motion be defined as

$$V_{\text{tan}} = r \omega,$$

representing the tangential velocity at the wing centre due to the turn rate ω . Under this approximation, the pulled (inner) wing tip experiences a tangential velocity of approximately $0.84 V_{\text{tan}}$, while the released (outer) wing tip experiences approximately $1.16 V_{\text{tan}}$. This asymmetry introduces a

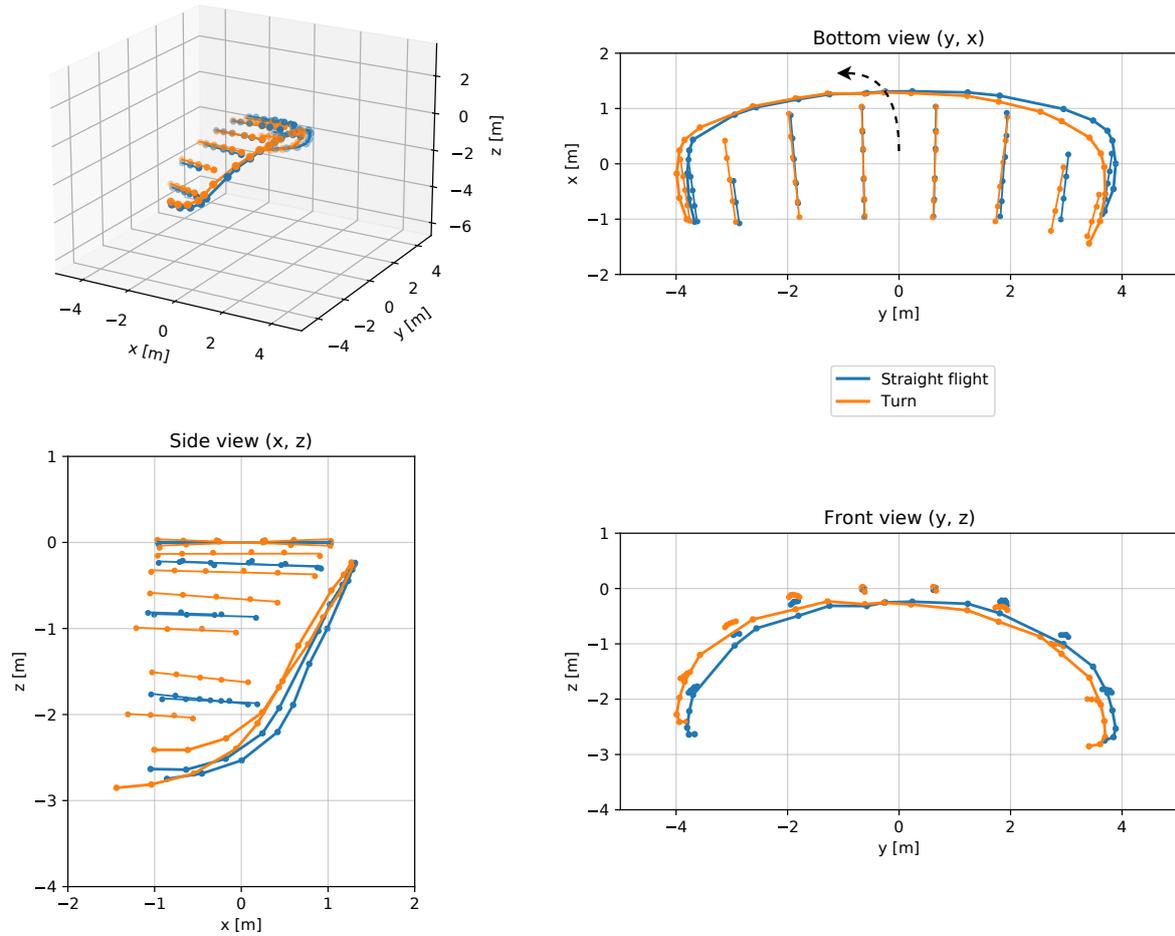


Figure 9.9: Shape change during a right-turn manoeuvre in the wing reference frame.

Table 9.5: In- and outputs for straight flight and the right turn shown in Fig. 9.9.

State	u_{dp} (-)	u_s (-)	b_{UWB} (m)	V_a (ms^{-1})	α_{fl} ($^\circ$)	β_s ($^\circ$)	F_T/g (kg)	$\epsilon_{p,adj,avg}$ (cm)	$\epsilon_{p,adj,max}$ (cm)
Straight flight	0.3937	0.0019	7.30	10.52	13.1	-7	102.81	2.91	4.19
Right turn	0.3937	-0.1869	7.25	15.18	13.1	-3	400.08	2.94	4.29

spanwise variation in V_a . However, as this effect is not directly measured and is further complicated by the magnitude and direction of the wind field, it is not taken into account in the present analysis.

In conclusion, the higher local angle of attack α_{local} on the pulled side results in an increase in drag, while the released side experiences reduced drag due to a lower local angle of attack, as depicted in Fig. 9.8. This differential drag causes the kite to yaw, thereby validating claim (i).

Redistribution of effective aerodynamic area

The front-view shear angle $|\Delta\xi|$ captures the vertical displacement of the tips during the manoeuvre. In the wing reference frame, the pulled tip moves upward, while the released tip moves downward, which may initially seem counter-intuitive for a turning manoeuvre. This displacement in the z -direction influences the aerodynamic area. As a result, the released side gains effective area, while the pulled side loses it. Although this is contrary to the expected behaviour, where the pulled tip moves closer to the KCU and the released tip moves farther away, this observation is a result of analysing the deformation in the wing reference frame. Despite the apparent contradiction, this displacement contributes to the roll moment as a result of the differential lift, which is consistent with theoretical predictions.

Asymmetric out-of-plane deformation

The deformation angles measured during the turn are:

$$|\Delta\zeta| = 4.92^\circ, \quad |\Delta\xi| = 4.39^\circ,$$

with uncertainties of $\pm 0.93^\circ$ and $\pm 0.98^\circ$, respectively, based on the average accuracy during the manoeuvres. The bottom-view shear angle $|\Delta\zeta|$ shows that the pulled tip moves upstream, while the released tip moves downstream due to slackened tension. Although the expected forward motion of the pulled tip is not fully supported by the data, this deviation can be attributed to the alignment of the data in the wing reference frame.

Comparison between wing and camera reference frames

When the same geometries are inspected in the camera reference frame (Fig. 9.10), a complementary pattern emerges. The camera view, rather than isolating the shape, preserves the global motion, most notably the combined roll–yaw rotation of the wing. This provides a clearer understanding of the global motion caused by steering, which was not fully captured in the wing reference frame where the wing is aligned.

In the camera reference frame, the pulled tip moves closer to the KCU and exhibits a noticeably larger projected area, while the released tip moves away and loses projected area. This behaviour is a direct result of the roll moment described earlier. This roll motion, induced by the shape change, is a consequence of the force imbalance between the two tips. This behaviour provides further validation for claim (ii), adding a deeper understanding of the underlying dynamics. Note that the results should be interpreted relative to straight flight, and not in absolute results, as the left tip appears closer to the KCU initially, solely because of the Γ -bar position.

The bottom and front views show that the entire wing rolls and yaws during the manoeuvre, which aligns with the predictions of the claims. In this frame, it is furthermore concluded that the pulled tip does move upstream, whereas the released tip moves downstream, validating claim (iii). This observation reinforces the idea that the kite does indeed roll and yaw, but it also introduces some complexity in defining ‘shape deformations,’ as this depends on the reference frame used for analysis.

The differences between the two reference frames highlight a crucial distinction:

- The wing reference frame reveals *pure shape deformation* by filtering out rigid-body motion. It quantifies local geometric deformation of the leading edge, canopy, and struts.
- The camera reference frame captures *global orientation changes* that also affect aerodynamic properties (roll, yaw and pitch of the wing).

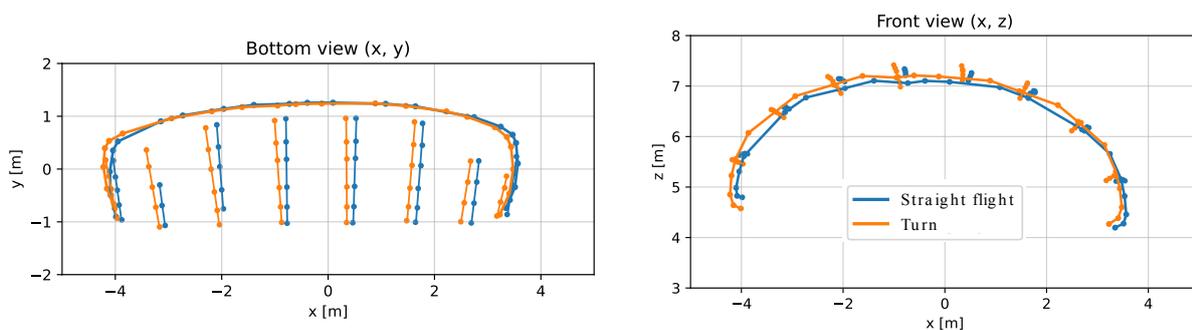


Figure 9.10: Bottom and front view of the right-turn geometry visualized in the camera reference frame.

Interpreting steering behaviour requires considering both components simultaneously. The local geometric deformation of the wing leads to force imbalances, which then result in the global reorientation of the kite.

Summary of relevant findings

The asymmetric twist distribution (claim (i)), redistribution of projected aerodynamic area (claim (ii)), and streamwise tip asymmetry (claim (iii)) generate the coupled roll–yaw moment that enables turning. The overall wing span decreases only minimally, which is consistent with localized deformation rather than a global contraction of the canopy. A summary of the quantitative changes is presented in Table 9.6.

Table 9.6: Summary of findings when going from straight flight to right turn (powered). Uncertainties are based on the average photogrammetry and UWB accuracies, found in Chapter 4 and Chapter 5 respectively.

Changing variable	Symbol	Value	Uncertainty
Steering constant	Δu_s	-0.1888	Unknown
Bottom shear	$ \Delta \zeta $	+4.92°	$\pm 0.93^\circ$
Front shear	$ \Delta \xi $	+4.39°	$\pm 0.98^\circ$
Span	Δb	-0.68%	$\pm 0.60\%$
Strut TE distance	$\Delta \bar{L}_n$	-1.60%	ND

A key outcome of this analysis is the strong dependence of the observed deformation on the chosen reference frame. The wing reference frame isolates pure shape change by removing global roll, yaw and pitch, but in doing so it obscures the global motions that are fundamental to turning. Conversely, the camera reference frame preserves these global motions, revealing changes in projected area and upstream/downstream displacement that are essential for understanding the aerodynamic consequences of steering. The combined interpretation shows that steering-induced deformation consists of both true geometric shape changes and global orientation changes. Only by analysing both reference frames together can the full turning mechanism be correctly understood.

The qualitative behaviour observed across all analysed cases is expected to be representative of LEI kites in general, as the trends measured during the selected manoeuvres are consistent with hypotheses reported in earlier studies. In contrast, the quantitative deformation levels (i.e. the absolute magnitudes) are likely specific to the TU Delft V3 kite, due to their strong dependence on the wing geometry and bridle configuration [142]. Applying the same experimental methodology to RAM-air and single-skin kites would enable a broader comparison of deformation mechanisms and help to further elucidate the structural role of inflatable components.

Structural model validation 10

One of the main aims of this research was to provide validation data/validation cases for aero-structural models. As the deformations validate the structural side of the model, this chapter discusses comparison with a structural finite element model (FEM). This chapter first introduces the FEM in Section 10.1. Afterwards, an experiment is constructed which produces validation cases for this model, the experimental setup and these test cases are discussed in Section 10.2. With the cases established, a method of validating the model is described in Section 10.3, based on shape correlation between experiment and model. Lastly, the results of this comparison is provided in Section 10.4

10.1 Structural finite element model

Currently, a structural model is being improved by P. Roeleveld [142], which can be used to model the deformations of the kite when subjected to external force inputs. This FEM of the TU Delft V3 kite was developed to predict how the kite deforms under external loads such as gravity, steering inputs, and additional point forces. The aim of the model is to capture the main structural mechanisms of the kite while being computationally efficient, such that it can be coupled to an aerodynamic model. The composition of the kite in this FEM is graphically depicted in Fig. 10.1.

The kite is represented using finite elements. The leading edge and struts are modelled as beam elements, which approximate the bending and torsional stiffness generated by the internal pressure inside the tubes. Instead of relying on traditional material properties, the stiffness of these beams is derived from experimental characterisations of inflatable structures as defined by Breukels [17]. This allows the model to reflect the non-linear behaviour typical of pressurised tubes, whose effective stiffness changes with deformation.

The canopy is discretised into quadrilateral sections. These are connected using non-compressive spring elements that largely prevent in-plane stretching but allow the surface to deform globally. This provides a simple but effective way to maintain the aerodynamic shape of the kite while still allowing the overall structure to flex in response to loads. The bridle system is also modelled as a

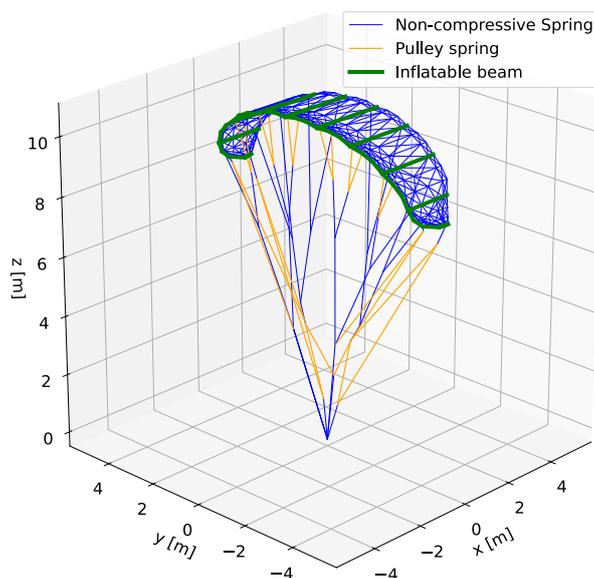


Figure 10.1: Kite composition in the FEM [142].

network of non-compressive springs, including the pulley mechanisms. These elements represent how forces are transferred from the canopy and beams to the tether connections, and play a large role in the structure’s shape. Because the kite undergoes large displacements and rotations, the model needs to resolve geometric non-linearity. This means that the element orientations and load paths change as the structure deforms, and the equilibrium shape must be found by iteratively updating the deformed geometry. The result is a static prediction of the kite’s shape under any given set of applied loads.

10.2 The hanging kite

To validate the aero-structural model under controlled, repeatable conditions, aerodynamic forces must be removed from the experiment. For this purpose, a fully static configuration—referred to as the hanging kite—was designed. In this setup, the kite is suspended upside down from a ceiling, such that gravity is the only external load, counterbalanced through the bridle system at a single suspension point. This allows direct comparison between measured deformations and the deformation model without the complexity of aerodynamic loading.

As no sufficiently tall indoor space was available, the bridle lines were shortened to achieve the required hanging geometry. The resulting bridle configuration, shown in Appendix G, caused the kite to have a slightly unconventional shape. In particular, the shortened back lines pulled the wingtips inward at the trailing edge, producing a pronounced positive twist. This altered bridle layout was simply and corrected in the FEM model, ensuring correspondence between experiment and model. Moreover, the revised geometry increased the loads on the bar tips, leading to measurable bar bending. This bending was identified and corrected during post-processing, as discussed in Section 8.2.1. A photograph of the experimental setup is shown in Fig. 10.2.

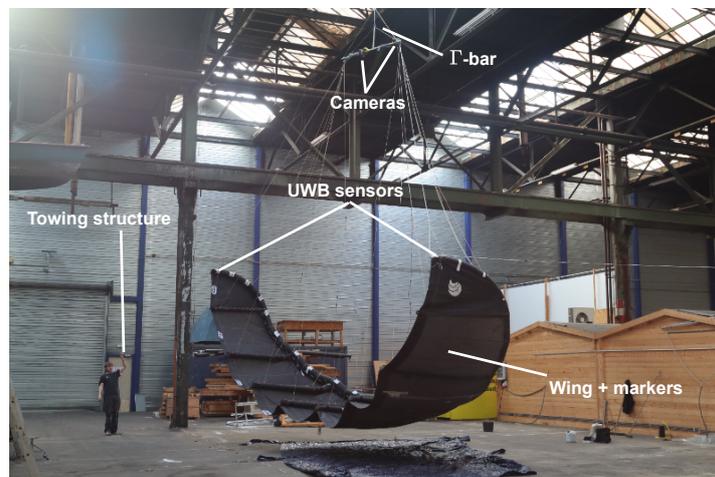


Figure 10.2: Hanging-kite experimental setup used for static deformation measurements.

The kite geometry was recorded using the photogrammetry system, with UWB-based sensor fusion applied to refine the stereoscopic calibration during post-processing, similar to the flight test discussed previously. To vary the load while maintaining accessibility, the kite was connected to a simple towing line routed over a horizontal bar. By adjusting this line, the kite could be raised or lowered, allowing variation of the LE and strut pressure as well as application of additional loads. Once positioned, the line was fixed, ensuring a static equilibrium corresponding to the “converged” state that the structural model aims to reproduce.

Test case definition

A set of symmetric test cases was defined to evaluate the model. Symmetry was preferred because the kite is free to rotate around the towing line, and asymmetric loads would introduce uncontrolled

rotations. Three experimental variables were considered:

- (i) **Inflation pressure:** The internal inflation pressure influences the overall stiffness of the structure and therefore its susceptibility to deformation. Two pressure levels were tested to assess the pressure sensitivity of the deformation model.
- (ii) **Point load:** A single symmetric point load was added by placing a mass at the central segment of the LE. Due to available weights, these took slightly unconventional values.
- (iii) **Tip load:** A symmetric outward pull on the wingtips was applied. One line was fixed at the tip height, while the other passed through a load cell to maintain a constant pulling force. Small misalignments are neglected under the small-angle approximation.

A graphical explanation of the load cases considered is shown in Fig. 10.3, showing both the point load case F_{PL} and the tip load case F_{TL}

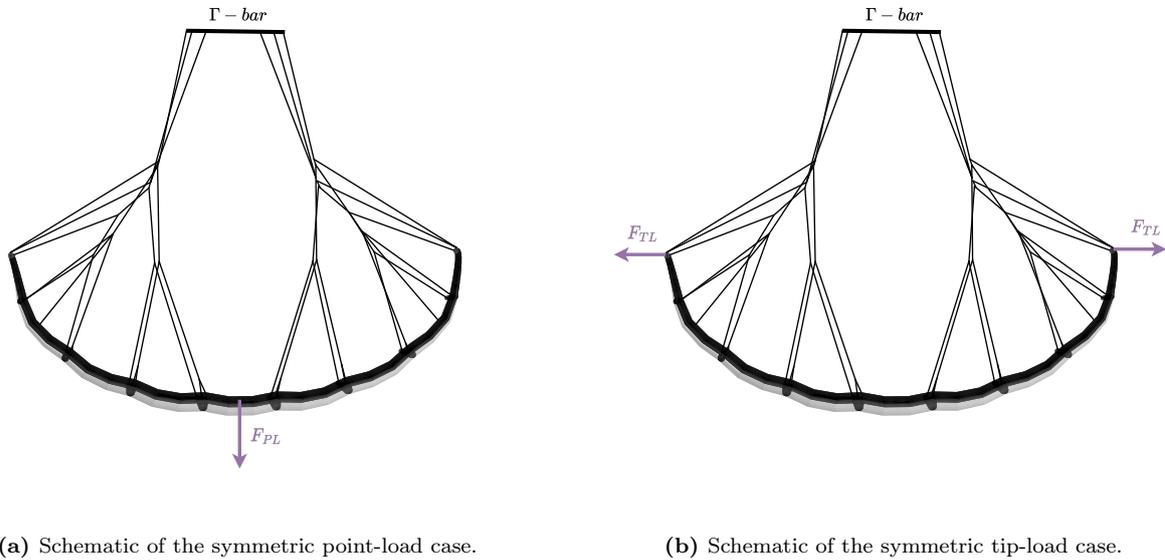


Figure 10.3: Overview of the load cases applied in the static hanging-kite experiment.

Summary of test cases

All test cases are summarised in Table 10.1. Two pressure levels were combined with two point loads and two tip loads. For each pressure, a static unloaded case (S) was included as a baseline.

The table also lists the stereoscopic calibration corrections required to compensate for bar bending during the static measurements. As expected, larger loads require larger corrections—particularly in yaw—which provides clear evidence that the bar bends under load. Interestingly, these increased

Table 10.1: Combined overview of static test cases with angle adjustments and applied loads.

Test case	p (mbar)	F_{PL}/g (kg)	F_{TL}/g (kg)	$\Delta\psi$	$\Delta\phi$	$\Delta\theta$	$\epsilon_{p,adj,avg}$ (cm)	$\epsilon_{p,adj,max}$ (cm)
P1 + S	150	0	0	0.772	0.030	0.202	2.36	4.13
P1 + PL1	150	9.7	0	1.210	-0.530	0.442	1.30	2.44
P1 + PL2	150	25.2	0	1.682	-1.040	0.448	3.48	7.10
P1 + TL1	150	0	2	1.342	-0.440	-0.126	2.27	5.11
P1 + TL2	150	0	5	1.404	-0.220	-0.248	1.61	4.21
P2 + S	250	0	0	0.982	-0.554	-0.184	1.90	4.00
P2 + PL1	250	9.7	0	1.142	-0.788	0	1.88	3.17
P2 + PL2	250	25.2	0	1.500	-1.140	0.308	2.42	3.99
P2 + TL1	250	0	2	0.926	-0.462	-0.110	1.56	3.65
P2 + TL2	250	0	5	1.070	-0.432	-0.272	2.48	6.01

corrections do not necessarily correspond to reduced reconstruction accuracy. This indicates that the correction procedure can successfully adjust the calibration retrospectively without degrading the accuracy levels established earlier (see Fig. 4.18). The average errors here are well below 5 cm, while the maximum errors are well below 10 cm.

10.3 FEM validation criterion

To validate the finite-element model (FEM) against the experimentally reconstructed deformation, a quantitative comparison metric is required. Although both the FEM and the photogrammetry describe the kite geometry along the same structural elements (leading edge and struts), the node locations do not coincide one-to-one. A direct pointwise comparison is therefore not feasible. Instead, a set of geometrically meaningful distances is extracted from both datasets to represent the overall deformation state. These distances are collected in the vector:

$$\Phi_n = \left[\underbrace{L_a \ L_b \ L_c \ L_d \ L_e \ L_f \ L_g \ L_h \ L_i}_{\text{Billowing segments}} \ \underbrace{b}_{\text{Span}} \ \underbrace{L_{CS-T,L} \ L_{CS-T,R}}_{\text{Centre-strut to tip distances}} \right]^T.$$

All distances are illustrated in Fig. 9.2. The first nine entries correspond to spanwise strut-to-strut distances along the trailing edge and capture local canopy curvature due to billowing. These segments primarily encode deformation in the y_w - and z_w -directions. The span b is included as a global measure of lateral contraction or expansion in the y_w -direction.

To incorporate deformation components in the x_w direction, the distances between the leading-edge point aligned with the centre strut and the corresponding tip are included on both sides of the wing. These quantities, $L_{CS-T,L}$ and $L_{CS-T,R}$, capture deformation in the x_w -direction and are sensitive to combined effects of twist, pitch, and longitudinal stretching. Together, the elements of Φ_n provide a compact yet sufficiently descriptive representation of the three-dimensional deformation state, while remaining independent of exact node correspondence.

The agreement between the FEM prediction and the experimental reconstruction is quantified using the mean absolute deviation

$$\epsilon_{\text{avg,exp,mod}} = \frac{1}{N} \sum_{i=1}^N |\Phi_{\text{exp},i} - \Phi_{\text{mod},i}|,$$

where the error is expressed in metres. This metric has a direct physical interpretation: it represents the average absolute geometric mismatch between the model and experiment over all selected deformation measures.

In addition to this absolute error, a dimensionless shape correlation coefficient SC is introduced to express the relative agreement between the two deformation patterns:

$$SC = 1 - \frac{\epsilon_{\text{avg,exp,mod}}(\%) }{100}.$$

Here, $SC \approx 1$ indicates a high correspondence between the FEM-predicted and experimentally observed deformation shapes, whereas values approaching zero indicate poor agreement. While SC does not capture localised discrepancies or phase shifts in deformation, it provides a convenient scalar measure for comparing performance across different load cases.

10.4 Results

Table 10.2 summarises the shape correlation coefficients and mean absolute errors for all ten static load cases. Three-dimensional visualisations of the corresponding experimental geometries are provided in Appendix H. For each case, the FEM deformation was evaluated using the validation procedure defined in Section 10.3.

Table 10.2: Shape correlation coefficients and mean absolute distance errors for all static load cases.

Test case	SC	$\epsilon_{\text{avg,exp,mod}}$ (m)
P1 + S	0.89	0.17
P1 + PL1	0.80	0.38
P1 + PL2	0.82	0.37
P1 + TL1	0.94	0.14
P1 + TL2	0.95	0.13
P2 + S	0.92	0.13
P2 + PL1	0.85	0.32
P2 + PL2	0.87	0.30
P2 + TL1	0.95	0.12
P2 + TL2	0.94	0.15
Average	0.89	0.22

Overall, a clear correspondence between the experimental reconstructions and the FEM predictions is observed across the majority of test cases. The shape correlation coefficients range between $SC = 0.80$ and $SC = 0.95$, with an average value of $SC = 0.89$. This indicates that the FEM captures the dominant deformation patterns reasonably well, particularly for load cases involving distributed loading along the structure.

The mean absolute errors remain non-negligible, with an average of $\epsilon_{\text{avg,exp,mod}} = 0.22$ m. This magnitude should be interpreted in the context of the experiment: the error metric aggregates discrepancies from multiple deformation measures and is therefore sensitive to both experimental uncertainty and modelling assumptions. Contributions arise from photogrammetry reconstruction errors, marker placement offsets, simplifications in the FEM material model, and the idealisation of load application.

The largest discrepancies occur in the point-load test cases (PL1 and PL2), where the average errors exceed 0.30 m and the shape correlation decreases noticeably. These cases involve highly localised and nonlinear deformation modes that are difficult to resolve accurately with the current FEM. In contrast, the static and tip-load cases (TL1 and TL2) exhibit consistently high correlations ($SC \geq 0.94$) and lower absolute errors, indicating that the model performs best when the deformation is dominated by global structural response rather than localised effects. The tip loads applied symmetrically stretches the entire wing by relatively the same factor, whereas the the point load substantially twists the centre struts, leaving other regions on the wing relatively unchanged. Despite the relatively large absolute errors in some cases, it is important to note that the FEM correctly predicts deformation trends over span changes of several metres across different load cases. This suggests that the dominant stiffness characteristics and load paths are captured adequately, even if local accuracy remains limited.

In summary, the FEM shows a clear and consistent agreement with the experimentally observed deformations for the majority of the static load cases, particularly those characterised by distributed loading. For highly localised point loads, the agreement deteriorates; nevertheless, the overall deformation trends and global structural behaviour are captured well. Importantly, the devised photogrammetric measurement method—originally developed for in-flight deformation quantification—was successfully applied to controlled static tests and enabled a direct comparison with the FEM results. Together, these findings indicate that the FEM provides a reliable first-order representation of the kite’s structural response, while also highlighting areas where further refinement, especially in modelling localised deformations, would be required to improve quantitative predictive accuracy.

Conclusions 11

This thesis set out to quantify the large-scale deformations of a leading-edge inflatable (LEI) kite used in AWE through in-flight experimental measurements. To achieve this, a measurement framework was developed that integrates stereoscopic photogrammetry with UWB ranging, enabling the continuous reconstruction of the kite’s three-dimensional shape during flight. The entire Python code is available online¹. The system was deployed on Kitepower B.V.’s AWES in Bangor Erris, Ireland, and validated further using controlled static experiments. The flight test data is also available online². Together, these campaigns yielded a rich dataset of structural deformations under realistic loading conditions. The results are concluded below based on the main research questions Section 3.2, which have been answered throughout the report. In addressing these research questions, this work directly fills the gaps identified in Chapter 2: the absence of established guidelines for measuring kite deformations, the lack of quantitative in-flight deformation datasets, and the shortage of experimental data suitable for validating aero-structural models. The measurement framework and results presented here conclude exactly these missing elements discussed in the report.

Experimental methods for in-flight deformation measurements (RQ1)

The developed measurement framework demonstrates that stereoscopic photogrammetry, supported by UWB span measurements, provides a practical, accurate, and reproducible method for measuring soft-wing deformations during outdoor flight. The dual-camera configuration with cross-shaped markers enabled reliable detection across a wide range of viewing angles, while the calibration pipeline produced stable intrinsic and extrinsic parameters. UWB fusion improved robustness by correcting small yaw and scale changes resulting from bar bending, ensuring centimetre-scale geometric accuracy throughout long-duration flights. When flown on an automated pattern at the Bangor Erris test site, the system also proved reproducible under naturally varying wind conditions, demonstrating reliable operation in realistic outdoor environments. Minor in-plane inaccuracies were observed during flight, mainly caused by challenging marker detection under strongly varying lighting conditions, but these did not compromise the overall reconstruction quality.

These findings were considered in light of alternative measurement techniques reviewed in the literature. While single-camera photogrammetry or LiDAR could also be used, both involve limitations—reduced depth accuracy for monocular systems and high cost and fragility for LiDAR. Techniques such as RTK-GPS or CDGPS show potential for span measurements and offer similar accuracy to UWB. Nevertheless, the photogrammetry–UWB combination offers a favourable balance of accuracy, intrusiveness, scalability, and cost.

Furthermore, the system proved adaptable: markers can be applied to different kite geometries, the camera baseline may be adjusted to match desired depth accuracy, and the bar can be mounted in various bridle configurations. Although the UWB modules used in this work were range-limited (≈ 25 m), newer technology or modified radio settings could easily extend this capability.

In-flight deformation behaviour (RQ2)

The in-flight results reveal clear and repeatable deformation modes characteristic of LEI kite operation. Across the flight envelope, the system successfully captured billow development, twist variation, changes in pitch angle associated with depower settings, and asymmetric tip motions during steering manoeuvres. These deformations were measured continuously, enabling detailed quantification of

¹https://github.com/pimjhaanen/photogrammetry_thesis

²<https://github.com/awegroup/Flightdata09102025>

both global and local shape changes relative to the undeformed CAD model. While the resulting deformation snapshots were found to be sufficiently accurate, the automated photogrammetry pipeline remains susceptible to errors. In particular, temporal consistency was limited, leading to increased noise and mismatches in the time-resolved deformation signals. This was in part caused by the bending and twist of the Γ -bar in flight, together with falsely identified markers, differential lighting, and outlier removal. Further refinement of the automation and tracking procedures would therefore improve the robustness of time-stepped analyses.

During depower inputs, the kite exhibited predictable structural responses. Billowing remained visible throughout flight, and twist distribution varied systematically. Depowering was found to reduce the pitch angle by an orientational change of the wing, confirming the expected reduction in aerodynamic angle of attack. Depowering also increased the anhedral angle, producing a more curved planform and shortening the wingspan, though the decrease was slightly lower than predicted by earlier simplified models. Average billowing decreased, consistent with the observed decrease in tether force.

Steering a LEI kite induces asymmetric deformations that cause both local shape changes and global reorientation, similar to depowering. In the wing reference frame, these deformations result in changes to the effective aerodynamic area, generating a roll moment. Furthermore, this results in a yaw moment due to a differential drag and lift loading on each tip, combined with the tips relative position. However, when viewed in the camera reference frame, the global roll and yaw of the kite become apparent. The combination of these local shape deformations and the resulting global motion explains the coupled roll–yaw behaviour observed during turning. These findings highlight the complex interplay between shape changes and global orientation shifts in the kite’s turning mechanism.

Comparison with the undeformed CAD geometry reveals several important differences. The real wing exhibited more billowing than the CAD model, causing the tips to tilt inward and producing a more bend planform. Twist angles were also positive in flight, indicating a larger local angle of attack. These differences have clear aerodynamic implications and emphasise the importance of using experimentally reconstructed shapes when validating aero–structural models.

Deformation model validation (RQ3)

The measured deformations are formatted such that these are well suited to support validation of aero–structural and fluid–structure interaction (FSI) models. The static hanging-kite tests provide controlled deformation cases that can be compared directly with finite-element simulations, particularly because the reconstructed leading-edge and strut geometries correspond to the structural points typically represented in deformation models. In addition, the in-flight dataset enables assessment of model behaviour under realistic operational loading conditions, including transient steering and depower manoeuvres. Overall, the FEM results showed a reasonable correspondence with the experimentally observed deformation trends.

Summary statement

In summary, this thesis demonstrates that accurate, synchronised, and continuous in-flight deformation measurements of LEI kites are achievable through a combined photogrammetry–UWB approach. The resulting dataset provides new insight into the structural behaviour of soft kites during operation and forms an essential basis for the development and validation of aero–structural models in AWE. Although the current implementation operates offline, the processing pipeline can be adapted for real-time operation using moderate hardware improvements, enabling deformation-aware control and early fault detection. As such, this work establishes a clear pathway from offline measurement capability to online improved safety, controllability, and energy yield in AWES.

Recommendations 12

This work establishes an experimental framework for in-flight deformation measurements of LEI kites and provides a validated dataset suitable for aero-structural model development. Several recommendations naturally arise from this research. They fall into two categories: (i) improvements to the present measurement setup and methodology (Section 12.1), and (ii) broader recommendations for future work enabled by the results of this thesis (Section 12.2).

12.1 Improvements to the current methodology

Bar design and structural stability: the primary limitation encountered in this study was the deformation of the Γ -bar under operational loads. Both bending and torsion were observed, particularly during high-tension flight phases and in the static hanging-kite experiment. Although these effects were corrected in post-processing through adjustments of the stereoscopic calibration matrices, the approach did not account for translations of the cameras relative to one another, which in reality also occur.

Future systems would benefit from a more rigid bar design—e.g. a carbon-fibre structure with higher stiffness-to-weight ratio. However, increasing rigidity alone may not eliminate compression forces associated with varying front-line tension. A practical solution is to allow one of the bridle clamps to translate and rotate along the bar. Such a sliding clamp would permit small relative motions without inducing structural deformation, provided that an end-stop mechanism prevents accidental detachment. This modification would substantially reduce the need for post-processing corrections, but bending might still occur as a result of aerodynamic drag. However, this effect is likely negligible when a more rigid bar is implemented.

Automation of correction factor identification: the manually determined correction factors used for the time-resolved visualisations of wing shape could be automated. Although not essential for the shape comparisons performed in this work, automated correction identification would reduce manual workload and improve repeatability. An optimisation routine could, for example, identify the tip markers in each frame (e.g. the LE markers with highest vertical pixel coordinates) and enforce consistency with UWB span measurements while simultaneously minimising epipolar discrepancy. This is basically an atomization of the approach discussed in Section 8.2.1. While computationally more expensive and potentially less accurate than manual selection, such automation would be valuable for large datasets or real-time applications.

Marker design and detection robustness: marker detection proved challenging due to highly variable outdoor lighting conditions, leading to in-plane uncertainties and occasional false positives. Improving marker design is therefore highly recommended. Using a kite with a uniform single-colour canopy would allow high-contrast marker designs (e.g. black–yellow crosses) and simplify contour identification. Even more effective would be the use of infrared reflective markers combined with an infrared light source or IR-enabled cameras, ensuring that only markers appear in the IR spectrum and eliminating false detections entirely. Markers should also be securely adhered using permanent adhesive rather than temporary tape to prevent detachment during repeated in- and deflations.

Distinguishing between LE and strut markers based on background brightness was effective but imperfect under varying lighting. Future work may incorporate more robust classification strategies, such as colour-invariant filtering or geometric model fitting, to ensure consistent wireframe reconstruction of the wing.

Reduction of KCU mass: in the present experiments, the KCU was significantly oversized relative to the kite, with a mass approximately twice that of the combined wing and bridle system. This was a consequence of using an off-the-shelf KCU designed for substantially larger kites. The excessive mass resulted in a relatively slack tether, which reduced controllability, manifested by lower tether forces and less pronounced dynamic responses. For future experiments, it is therefore recommended to use a KCU with a mass more closely matched to the kite size, in order to ensure sufficient tether tension and representative flight dynamics.

Sensor synchronisation: although synchronisation across the KCU, pitot tube, angular vanes, and UWB units was successfully achieved using UTC timestamps, as well as the two GoPro cameras with one another, these cameras could not be perfectly synchronised to UTC due to variable LED trigger delays. Future studies could therefore benefit from compact cameras capable of external time-source synchronisation. Such an approach would allow frame-perfect alignment across all sensor systems without requiring manual intervention.

Expanding the sensor system: direct measurement of aerodynamic angles, particularly wing angle of attack, would significantly enhance the interpretation of kite deformation. The angular vanes mounted on the front bridles provided limited information on wing angle of attack but were too intrusive to be positioned on the wing. Altering the setup such that wing angle of attack is measured would provide significantly more insights on the behaviour, particularly during depowering.

On top of this, angular measurements calculated through the point cloud (e.g. pitch) often lacked accuracy. Future setups could therefore incorporate lightweight inertial or angular sensors on selected structural elements such as struts or wingtips, providing direct measurements of pitch and twist. Additionally, extending sensing to the canopy could provide a more complete three-dimensional representation of billowing behaviour.

Improving the photogrammetry algorithm: while the individual deformation snapshots were found to be highly accurate, the automated reconstruction of wing deformations can be further improved. The current pipeline is sensitive to challenges in marker detection, including non-uniform lighting conditions, occasional false detections, and the additional corrections required to compensate for bar bending. As a result, the fully automated output is at times noisy and susceptible to errors.

Although addressing these limitations fell outside the scope of the present thesis, future work could focus on refining marker detection robustness and improving outlier rejection strategies. One promising approach would be to use the undeformed CAD geometry as a reference model, enabling outliers to be identified and removed based on their deviation from a physically plausible wing shape.

12.2 Future research directions

Use of reconstructed shapes in aerodynamic analysis: the reconstructed in-flight shapes differ meaningfully from the idealised CAD model traditionally used in aerodynamic simulations. Incorporating experimentally derived billow profiles, twist distributions, and tip sweep into CFD or aero-structural models would yield more realistic aerodynamic predictions and enable improved kite designs. Adopting such corrected shapes may significantly influence load estimation, steering dynamics, and energy yield predictions.

Real-time shape tracking: the demonstrated measurement framework opens the possibility of real-time structural monitoring. Although the current implementation is limited by video quality, marker design, and processing time, modest computing upgrades or reduced video resolution could enable (near) real-time deformation reconstruction. This would allow continuous monitoring for early fault detection and pave the way for deformation-aware control strategies, enhancing safety and performance in AWES.

List of requirements A

This appendix contains a list of requirements used throughout this research. To ensure successful execution, a structured set of requirements is defined across the experimental setup, instrumentation, data processing, validation, and practical constraints. Together, these requirements provide a framework for an accurate, reliable, and reproducible measurement system under variable outdoor wind conditions. While formulated for the TU Delft V3 kite (see Chapter 1), they can be adapted to other LEI kites with different geometries or sizes. Not all requirements apply equally to every technique, as some rely on fundamentally different measurement principles. The complete list of requirements is provided below.

A.1 Experimental requirements

ID	Requirement	Accomplished
REQ-EXP-01	The experimental setup should achieve sufficient accuracy to generate reliable data.	N
REQ-EXP-01a	Distance or position data must have a minimum deformation measurement accuracy lower than 10 cm.	Y
REQ-EXP-01b	The depth error shall be no more than 5 cm.	Y
REQ-EXP-01c	Angular deformation measurements must have a minimum accuracy of 1° .	N
REQ-EXP-02	The test methodology must ensure that key parameters can be controlled despite varying wind conditions, enabling reproducible tests.	Y
REQ-EXP-03	The measurement system must collect data at a frequency of at least 4 Hz to capture vibrational deformations, at least twice the highest frequency present [40]; a higher rate is preferred.	Y
REQ-EXP-04	The setup should remain functional under expected wind speeds of $1 - 35 \text{ ms}^{-1}$.	Y
REQ-EXP-05	The measurement system should be adaptable for different sizes of LEI kites.	Y
REQ-EXP-06	The experiments must present real-life flight conditions, resembling manoeuvres present during the pumping cycle.	Y
REQ-EXP-06a	Deformations should be measured under different depower settings u_{dp} to analyse their effect on wing shape.	Y
REQ-EXP-06b	The experiment should capture deformations for different steering inputs u_s to analyse their effect on wing shape.	Y
REQ-EXP-07	The test methodology should minimise external disturbances.	N

A.2 Instrumental requirements

ID	Requirement	Accomplished
REQ-INS-01	Sensors must have a minimum resolution of centimetre level to detect small-scale deformations.	Y
REQ-INS-02	The instrumentation should be mountable and removable without effort.	Y
REQ-INS-03	All sensors must be synchronised in time.	Y
REQ-INS-04	The measurement setup should be minimally intrusive.	Y
REQ-INS-04a	The instrumentation must be present in positions which minimally interfere with flow behaviour around the kite.	Y
REQ-INS-04b	The instrumentation must weigh no more than 2.5 kg, which is equal to the mass of the wind turbine. This wind turbine can be removed, thus not adding any additional mass to the setup.	Y
REQ-INS-05	Sensors and equipment must be able to withstand impacts from potential kite crashes.	Y
REQ-INS-06	The instrumental setup must map the kite from below, for global shape mapping technologies.	Y
REQ-INS-06a	The instrumentation must have a range of at least 10 m.	Y
REQ-INS-06b	The instrumentation must have a horizontal field of view of at least 75°.	Y
REQ-INS-06c	The instrumentation must have a vertical field of view of at least 30°.	Y
REQ-INS-07	The sensors must be capable of determining both position or distance measurements for key nodes in space, if applicable.	Y
REQ-INS-07a	Measurement equipment must have a range of at least 10 m when measured from the kite.	Y
REQ-INS-07b	Measurement equipment must have a range of at least 400 m when referenced from the tether anchor point at the ground.	N/A
REQ-INS-07c	Measurement equipment must be omni-directional to allow tracking of nodes in all spatial directions.	Y
REQ-INS-08	The sensors must be capable of kinematic sensing to accurately capture the rotational motion and reconstruct 3D shape deformation of the kite.	N/A
REQ-INS-08a	The angular measurement accuracy must be at least 1°.	N/A
REQ-INS-08b	The measurement system must allow for full 3D reconstruction at the sensor location.	N/A
REQ-INS-08c	Sensors should maximise coverage to reduce reliance on interpolation between measurement points.	N/A
REQ-INS-09	Instrumentation should be present to measure the apparent wind velocity.	Y
REQ-INS-09a	The apparent wind speed should be determined on-board of the kite.	Y
REQ-INS-09b	The wind speed should have an uncertainty of no more than 1 ms ⁻¹ .	Y

ID	Requirement	Accomplished
REQ-INS-09c	The apparent wind speed sensor should be calibrated beforehand.	Y
REQ-INS-09d	The directional sensitivity of the apparent wind speed sensor should be determined and concluded no issue for the angles experienced between wing and sensor.	Y
REQ-INS-10	In case of sensor failure, the instrumentation should be rebootable within an hour.	Y
REQ-INS-11	Sensors must withstand environmental conditions.	Y
REQ-INS-11a	Sensors must remain functional within 10–30°C.	Y
REQ-INS-11b	Sensors must be protected from excessive moisture or precipitation.	Y
REQ-INS-11c	Sensors must withstand UV exposure from the sun.	Y
REQ-INS-11d	Sensor performance shall not degrade under different environmental conditions.	Y

A.3 Data processing and validation requirements

ID	Requirement	Accomplished
REQ-DPV-01	The measurement uncertainty must be limited and quantifiable.	N
REQ-DPV-01a	The measurement uncertainty must not exceed 1 decimetre.	Y
REQ-DPV-01b	The measurement uncertainty must not exceed 1° for angular measurements.	N
REQ-DPV-02	The data processing pipeline should incorporate a method to remove sensor noise.	Y
REQ-DPV-03	Measurement results must be comparable across different studies, sensors, experiments, and models.	Y
REQ-DPV-03a	Measurement results must be comparable with previous studies by ensuring the same deformation types are observed.	Y
REQ-DPV-03b	Measurement results must be comparable across different sensors used in the experiment, ensuring cross-sensor reliability.	Y
REQ-DPV-03c	Measurement results must be comparable between different experimental setups and conditions to ensure reproducibility.	Y
REQ-DPV-03d	Measurement results must be formatted and processed in a way that allows direct validation with FSI models.	Y
REQ-DPV-04	The experimental data should be scalable to different kite sizes.	Y
REQ-DPV-05	Wind conditions during experiments must be incorporated in the data acquisition.	Y
REQ-DPV-06	The data processing system must have a defined strategy for handling missing data points.	Y

ID	Requirement	Accomplished
REQ-DPV-07	Anomalous data should be flagged and reviewed manually before validation.	Y
REQ-DPV-08	Data must be kept synchronised to ensure accurate cross-sensor comparison or data fusion.	Y
REQ-DPV-09	Measurement results must be validated.	Y

A.4 Logistical and practical requirements

ID	Requirement	Accomplished
REQ-LP-01	The full experiment, including setup, execution, and data analysis must be feasible within 32 weeks.	Y
REQ-LP-02	The total budget for equipment and field tests must not exceed €2000.	Y
REQ-LP-03	The measurement system should be lightweight and easily deployable for field testing.	Y
REQ-LP-04	All experimental setups must comply with TU Delft and Kitepower safety guidelines.	Y
REQ-LP-05	One experimental test may last no more than one workday.	Y
REQ-LP-05a	The entire experimental setup must be deployable within 2 hours.	Y
REQ-LP-05b	The experiments, including different flight manoeuvres, must be doable within 4 hours.	Y
REQ-LP-05c	The entire experimental setup must be packable within 2 hours.	Y
REQ-LP-06	Data storage must allow saving data at the minimum sampling frequency for at least one test session.	Y
REQ-LP-07	Backup copies of the data must be stored in at least two separate locations to prevent data loss.	Y

Sensor trade-off B

This appendix provides the trade-off between different measurement techniques possible for kite deformation measurements. Based on insights from the literature review, candidate measurement techniques are evaluated against the requirements. To achieve this, the measurement techniques are divided into three primary measurement approaches, categorized based on their role in mapping kite deformations, each with their own set of specific requirements:

1. **Global shape mapping technologies:** These methods focus on capturing the entire shape of the kite, scanning from underneath the kite. This category consists mainly of visual measurement technologies, but also includes LiDAR scanning.
2. **Time-of-flight (ToF) and active ranging technologies** These techniques focus on position and distance tracking between specific nodes on the kite. They measure relative position or distance between structural components, providing spatial information such as strut-to-strut width or bridle line movement. These sensors can be mounted on the kite itself, the KCU, or the ground, depending on the instrumentation.
3. **Structural and kinematic sensing technologies:** These involve localized sensors embedded directly in the kite’s structure to measure strain, bending, and rotational motion. Technologies such as fiber optic strain sensors, Inertial Measurement Units (IMUs), and deformation sensors detect localized shape changes, allowing for kinematic analysis and interpolation-based 3D reconstruction of deformations.

Based on the insights from the literature review, a trade-off has been performed. The trade-off table can be found in Table B.1, whereas Table B.2 provides the legend. Scores range from -2 (poor) to 2 (outstanding). Techniques can be disqualified (DSQ) if a requirement is fundamentally unmet. The accompanying notes clarify the scope of each criterion and how scores were assigned. A discussion per technique is included in Chapter 2, below the evaluation is made based on the information gathered throughout the literature.

B.1 Global shape mapping technologies

The first category comprises global shape mapping techniques that observe the kite’s deformation from the KCU or from a position beneath the wing. Among camera-based methods, stereoscopic photogrammetry offers the highest accuracy—especially for out-of-plane motion. Single-camera photogrammetry is simpler but lacks depth perception for detailed reconstruction. Techniques such as 3D-GS and load cells are excluded due to insufficient accuracy, and structured-light 3D scanning—while precise at short range—is impractical at the required distances (REQ-EXP-01).

All candidate techniques exceed the minimum viable sampling frequency (REQ-EXP-03), ensuring that motion is captured effectively. However, intrusiveness and mass (REQ-EXP-04) vary significantly, considering not only sensor size but also associated wiring and data acquisition units. Sensors mounted farther from the wing cause less aerodynamic disturbance than those directly on the kite.

Given that kite crashes occur periodically, durability is a critical factor (REQ-INS-05). Techniques with fragile components, such as LiDAR, may require additional protective measures. For range requirements, global shape mapping instrumentation must function between 2 – 20 m, ensuring compatibility across kite sizes. Additionally, the field of view (FOV) is assessed with a margin to account for yawing and vibration effects [30] (REQ-INS-06).

Table B.1: Trade-off of different measurement techniques for mapping the entire kite, measuring node position/distance or structural/kinematic sensing technologies. The trade-off criteria are derived from the requirements.

* Accuracy is judged for the stated purpose and range. For IMUs, strain tapes and flex sensors, positional accuracy is not considered.

** Despite possibly being expensive, these units are readily available for use.

Instrumentation requirements													
Global Shape Mapping Technologies (from KCU)													
Measurement Technique	Accuracy (<1 dm)*	Sampling Frequency (>4 Hz)	Size/ Intrusiveness	Mass (<2.5 kg)	Can Withstand Impact	Range (>10 m)	Field of View (>75 deg H, >30 deg V)	Withstand Environmental Conditions	Performance Varying Environment	Cost/ Availability	Complexity	TOTAL	
(Single Camera) Photogrammetry	-1	2	2	2	2	1	2	2	-1	2**	1	14	
Stereoscopic Photogrammetry	1	2	2	2	2	2	2	2	0	2**	0	17	
Digital Image Correlation (DIC)	1	2	2	2	2	-1	2	0	-2	2**	-2	8	
Gaussian Splatting (3D-GS)	DSQ	2	2	2	2	-1	2	2	-1	2**	1	DSQ	
Structured Light 3D Scanning	-1	1	2	2	1	-2	1	0	DSQ	-1	-1	DSQ	
Depth Camera's	-1	1	2	1	0	0	0	0	0	0	2	5	
Time of Flight (ToF) Camera's	0	1	2	1	0	0	1	0	1	-1	1	6	
LiDAR	2	1	2	1	-2	2	2	1	2	-1	1	11	
Load cells in bridle system	DSQ	2	-1	0	2	N/A	N/A	2	2	0	-2	DSQ	
Time-of-Flight (ToF) and Active Ranging Technologies													
Measurement Technique Ground-Based : GB Kite-Based : KB	Accuracy (<1 dm)*	Sampling Frequency (>4 Hz)	Size/ Intrusiveness	Mass (<2.5 kg)	Can Withstand Impact	Kite Range (>10 m)	Ground Range (>60 m)	Omnidirectionality	Withstand Environmental Conditions	Performance Varying Environment	Cost/ Availability	Complexity	TOTAL
Radar (GB)	DSQ	2	2	1	0	N/A	2	-1	1	0	1	1	DSQ
Ultrasonic Distance Measuring (KB)	-1	2	1	1	0	-1	N/A	-2	-1	DSQ	1	2	DSQ
Global Positioning System (GPS, GB)	DSQ	2	2	2	1	N/A	2	2	2	-1	1	2	DSQ
Real-Time Kinematic GPS (RTK GPS, GB)	1	2	2	2	1	N/A	2	2	2	1	-1	1	15
Carrier-Phase Differential GPS (CDGPS, GB)	1	2	2	2	1	N/A	2	2	2	2	-1	1	16
Ultra-Wideband Modules (UWB, KB)	0	2	1	2	1	2	N/A	2	1	1	2**	2	16
Retro-reflective sensors (KB)	2	2	1	1	0	-1	N/A	DSQ	1	-2	1	2	DSQ
Diffuse sensors (KB)	2	2	1	1	0	-1	N/A	DSQ	-1	DSQ	1	2	DSQ
Thru-beams (KB)	2	2	1	1	0	2	N/A	DSQ	2	DSQ	1	1	DSQ
Draw-wire sensor (KB)	0	1	0	0	0	-2	N/A	0	0	-2	1	2	0
Linear potentiometer (KB)	2	2	0	0	0	DSQ	N/A	0	0	-2	1	1	DSQ
Structural and Kinematic Sensing Technologies													
Measurement Technique	Position Accuracy (<1 dm)*	Angular Accuracy (<1 deg)	Sampling Frequency (>4 Hz)	Size/ Intrusiveness	Mass (<2.5 kg)	Can Withstand Impact	Use for 3D Recon-struction	Spatial Coverage	Withstand Environmental Conditions	Performance Varying Environment	Cost/ Availability	Complexity	TOTAL
Inertial Measurement Units(XSENS Dot IMU's)	N/A	0	2	2	2	1	0	-2	2	0	2**	-1	8
Strain Tapes	N/A	0	2	-1	1	1	-1	-2	1	-1	1	-1	0
Flex Sensors	N/A	1	2	-1	1	0	-2	-1	1	-1	1	-1	0
Fibre Optic Shape Sensors (FOSS)	2	2	2	2	-1	0	2	2	2	2	DSQ	0	DSQ

Table B.2: Legend for Trade-Off table

Legend					
Outstanding: 2	Good: 1	Acceptable: 0	Bad: -1	Terrible: -2	Disqualified: DSQ

Environmental resilience (REQ-INS-11) encompasses two key aspects. First, the instrumentation must withstand exposure to environmental conditions such as rain, snow, hail, temperature variations, and UV radiation without failure. Sensors that are not designed for outdoor use or require excessive protection are less suitable. Second, performance stability in varying environments is evaluated. Some techniques, particularly vision-based methods, experience significant degradation in bad weather (e.g., rain and fog), even if the hardware itself remains functional. Structured light scanning, for example, is highly susceptible to these conditions and requires controlled environments, leading to its exclusion. Conversely, techniques like UWB and CDGPS are largely unaffected by environmental changes and maintain stable performance.

Cost and availability (REQ-LP-02) consider not only the price of instrumentation but also the cost of required software, data acquisition units, and any additional accessories needed for full functionality. Techniques that require purchasing expensive sensors or software receive lower ratings. If the equipment is already available for use, a score of 2 is assigned, whereas techniques that require acquisition receive a maximum score of 1. Complexity is assessed comprehensively, accounting for multiple factors that influence the practicality of the measurement system within the project timeline. This includes calibration difficulty, setup and tear down time (REQ-LP-05), wiring and cabling, ease of use, data acquisition requirements and data processing demands (REQ-LP-06). Techniques that require minimal wiring, easy calibration, and simple data processing are rated more favourably, while those with high setup complexity, intricate data handling, or extensive cabling score lower.

B.2 Time-of-Flight (ToF) and active ranging technologies

For time-of-flight and active ranging technologies, the range criteria differ for kite-based and ground-based systems. Kite-based systems require a 10 m range, while ground-based systems must function up to 60 m. Unlike global shape mapping, these systems must also be omnidirectional, as orientation between sensors or transmitter/receiver changes dynamically (REQ-INS-07).

Among these, radar and GPS-based methods lack the necessary accuracy, disqualifying them. Ultrasonic distance sensors, while offering slight angular coverage, are unsuitable due to motion of the kite. Similarly, retro-reflective sensors, diffuse sensors, and thru-beam systems require precise alignment, making them infeasible for airborne deformation tracking. One could argue that draw-wire sensors and linear potentiometers are less unidirectional, as their ends can be fixed to a moving source while still determining distance. However, the linear potentiometer is disqualified due to its low range, making it unsuitable for the required measurements. The draw-wire sensor remains under consideration, though its performance in longer-range applications is limited by potential line sag and drag effects.

B.3 Structural and kinematic sensing technologies

Lastly, there is a category of structural and kinematic sensing technologies, where angular accuracy is a key factor in their evaluation. These sensors are typically local measurement tools, meaning they do not inherently provide full 3D shape reconstruction but rather track specific deformations at discrete points. Therefore, an important consideration is their spatial coverage, which is how much deformation data they provide and how many assumptions must be made to interpolate the overall shape of the kite (REQ-INS-08). A key example is Inertial Measurement Units (IMUs), which offer high accuracy in roll and pitch measurements but suffer from significant drift in yaw and absolute positioning. While IMUs would be disqualified based on position accuracy alone, this is not the intended application, so they are retained in the evaluation. However, their ability to capture the kite's full shape remains limited.

Strain tapes and flex sensors, while capable of capturing local deformations, require extensive sensor networks to provide meaningful shape reconstruction. This means they demand a large number of sensors, complex cabling, and data acquisition systems, making their deployment impractical. Additionally, these sensors primarily measure 1D strain or bending at fixed locations, necessitating significant assumptions to extrapolate a 3D representation of the kite's shape. Even when multiple sensors are used, complex deformation patterns such as twisting and bending remain difficult to quantify with high accuracy. Fibre optic shape sensors (FOSS) perform exceptionally well, meeting nearly all requirements, including accuracy, spatial coverage, and reliability. However, the mass of the data acquisition unit can be a concern, and more critically, the cost of such a system is prohibitively high, especially when scaling up to a full-sized kite. Due to this, FOSS is excluded from consideration, despite its otherwise strong performance.

Fresnel zone calculation C

This appendix discusses the Fresnel zone calculation required for the UWB range test. The Fresnel zone was analysed as it constrains the effective transmission distance. The Fresnel zone is a series of ellipsoidal regions around the line of sight between a transmitter and receiver that strongly affect signal strength in wireless communication [133]. Maintaining a clear Fresnel zone is essential for minimising signal loss. For the test, only the largest radius (at the midpoint of the path) was considered, as this represents the minimum required ground clearance. The radius of the first Fresnel zone is given by [143]:

$$r_1 = \sqrt{\frac{\lambda \cdot d_1 \cdot d_2}{d_1 + d_2}},$$

where r_1 is the radius of the first Fresnel zone, λ is the wavelength of the transmitted signal, and d_1 and d_2 are the distances from the point of interest to the transmitter and receiver, respectively. A schematic of the Fresnel geometry is provided in Fig. C.1.

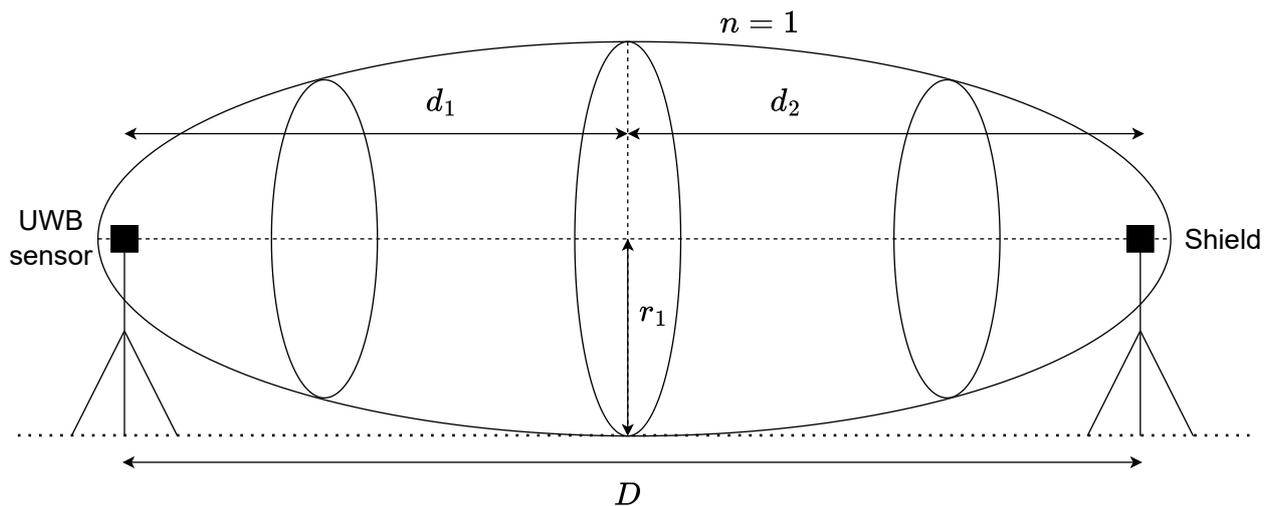


Figure C.1: First Fresnel zone geometry. The maximum radius at the midpoint defines the required ground clearance for maintaining unobstructed line-of-sight communication.

For the largest cross-sectional area of the zone, where the point of interest lies exactly halfway between transmitter and receiver ($d_1 = d_2 = \frac{D}{2}$), the expression reduces to:

$$r_1 = \sqrt{\frac{\lambda \cdot (\frac{D}{2})^2}{D}} = \sqrt{\frac{\lambda \cdot D}{4}}.$$

With a UWB carrier frequency of 6.5 GHz, the wavelength is $\lambda \approx 0.046$ m. For a total transmission distance $D = 60$ m (representing tether and bridle length during experiments as in REQ-INS-07b), the Fresnel radius was $r_1 = 0.83$ m. Applying a safety factor $SF = 2$ yields a minimum clearance of 1.66 m, similar to the unobstructed connection expected whilst experimenting. Additionally, the shield was therefore oriented vertically to maximise effective range.

Improved vane design D

This appendix discussed the updated angular vane design. Initial tests used off-the-shelf miniature vanes. Because these sat within the bluff-body wake of the casing, they experienced separated, unsteady flow that caused misalignment and vibration (flutter), which undermined measurement reliability. To address this, a stiffer and more stable custom vane was developed; the previous and updated designs are shown in Fig. D.1.

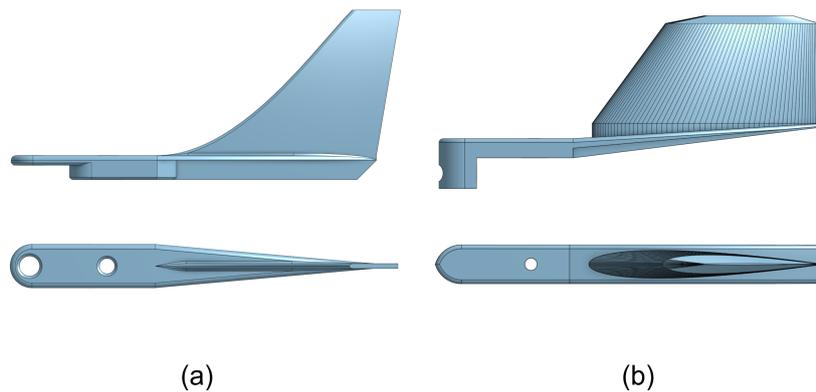


Figure D.1: Improved angular vane design (not to scale). (a) original flat vane located close to the casing; (b) redesigned vane with symmetric airfoil section and increased stand-off from the casing.

The redesign follows three principles:

1. **Symmetric, stable section.** The flat plate was replaced by a NACA 0012 airfoil. The “0012” indicates a symmetric profile with a maximum thickness of 12% of the chord. Symmetry minimizes bias about 0° and yields a small restoring aerodynamic moment in attached flow, improving linearity and reducing susceptibility to flutter compared with a thin flat plate.
2. **Reduced wake interference.** The vane was mounted farther from the casing on a slender support to move it out of the strongest part of the casing wake. This reduces local flow angularity and velocity deficits at the sensing location, thereby lowering systematic overshoot at larger inflow angles.
3. **Balanced, stiffer body.** The vane was 3D-printed with variable infill so that the mass center lies on the rotation axis (minimizing gravitational torque) while increasing torsional stiffness. The combination raises the onset speed for vibration and damps small-amplitude chatter.

Wind-tunnel tests confirmed markedly improved behaviour: the redesigned vane exhibited no visible flutter in the relevant speed range and showed smaller overshoot at high angles. A residual, weaker influence of the casing wake remains at the largest deflections, but overall the new configuration provides substantially more reliable angular measurements.

Custom printed parts E

This appendix discussed the design and reasoning behind the custom 3D-printed parts produced for this experiment. All mechanical interfaces are realised with custom 3D-printed components, shown in Fig. E.1. The parts were fabricated on a Bambu Lab X1 Carbon using PETG-HF with 35% gyroid infill, chosen to balance stiffness, impact resistance, and mass. The clamp designs ensure that no rotation or creep occurs under bridle loads, while still allowing rapid installation and removal when the lines are slack (REQ-LP-05, REQ-INS-02). The modular nature of the design means that the same hardware can be scaled to other wings by varying tube lengths to adjust the stereo baseline B (REQ-EXP-05). The different printed components of this experimental structure are described below:

1. **GoPro camera mounts.** Each mount employs a through-pin that penetrates the bar together with a high-strength clamping saddle, effectively eliminating any possibility of rotation or translation of the clamp. The camera casing is press-fitted into the holder and secured by a transverse bolt, while a recessed pocket underneath prevents yaw or roll under load. Additional sliders guide the bridles to move freely along the bar, avoiding interference. These precautions are particularly critical for the cameras, which are stereo-calibrated to this exact configuration. Any displacement or rotation—whether from bridle entanglement or high-impact events such as kite crashes—would compromise calibration. The mounts therefore incorporate multiple redundancy features to ensure positional stability throughout the experiments.
2. **Sensor box.** Mounted on the cross-member, this enclosure contains the UWB receiver (shield v1.5), Raspberry Pi, and a compact powerbank. An integrated LED window provides optical synchronisation with the camera system, discussed further in this chapter. The sensor box is rather large in size to allow for easy waterproofing using duct-tape.
3. **Pitot tube mount.** The mount positions the probe slightly in front and below the bar to reduce flow interference. The Pitot tube and the angular vanes are roughly aligned with the front lines and fixed tightly to avoid rotation, ensuring that in-flight angular misalignment remains within the reliable range established in Chapter 6 ($\epsilon_{ws} < 1 \text{ ms}^{-1}$ for $|\phi| \lesssim 35^\circ$).
4. **T-piece bar connector.** A printed junction rigidly couples the horizontal cross-member to the vertical leg, forming the Γ structure that provides torsional and bending stiffness.
5. **Vertical and horizontal bridle clamps.** Bridle lines are woven through shaped passages in the clamps. Under load, the line locks the clamp securely, making detachment impossible. With slack lines, the same design permits easy mounting and removal.

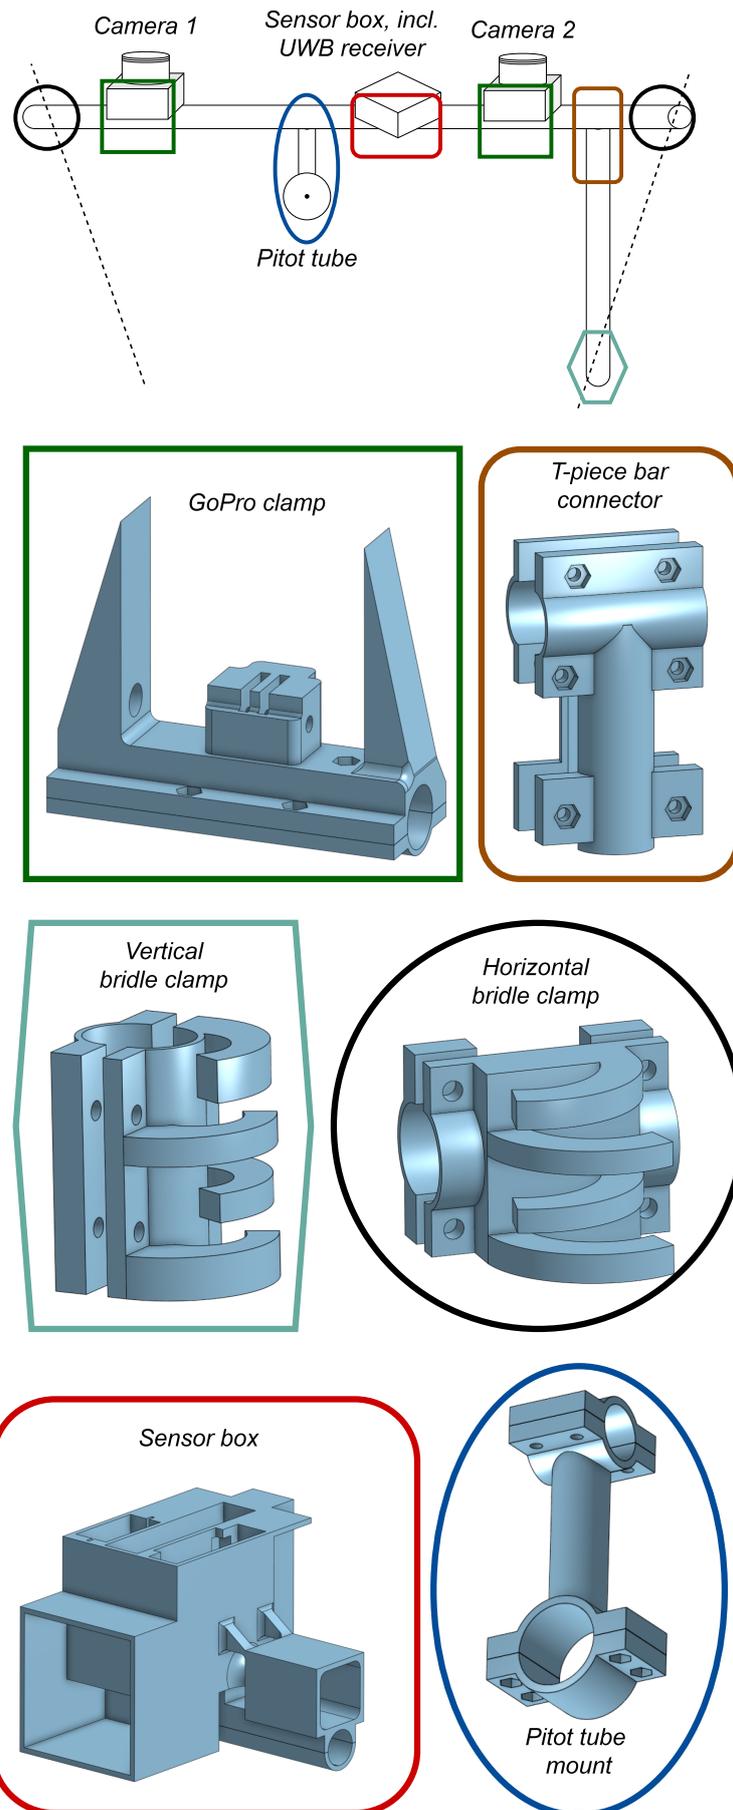


Figure E.1: Custom 3D-printed components used in the experimental setup. All parts were manufactured in PETG HF with 35% gyroid infill to balance stiffness, strength, and weight.

Bridle geometry flight tests F

This appendix documents the bridle geometry of the V3 kite configuration used during the flight tests. The complete set of measured line lengths is provided in Table F.1. These measurements were performed to ensure that the setup can be copied in either future experiments or structural models. One can observe all load-carrying connections are indeed positioned on the LE tube and the struts, and not on the canopy. The corresponding attachment locations on the wing are shown in Fig. F.1.

For clarity, the numbering conventions of the front and rear bridle systems are reproduced in Fig. F.2 and Fig. F.3. These schematics illustrate the structure of the bridle network and its subdivision into main lines and secondary segments. Black dots indicate knots whilst the white circles indicate pulleys.

Table F.1: Bridle length table used during the flight experiments.

Line ID	Length (mm)	Line ID	Length (mm)
Front Lines		Rear Lines	
A_{main}	3903	M	4856
A_{I}	3670	$B_{r,\text{main},e}$	70
A_{II}	3360	$B_{r,\text{main},1}$	610
$A_{\text{III}}/B_{r,\text{main},2}$	11420	$B_{r,\text{main},2}/A_{\text{III}}$	11420
A_1	1860	$B_{r,I}$	2333
A_2	1798	$B_{r,II}$	2045
A_3	1791	$B_{r,5}$	12168
A_4	2877	$b_{r,1}$	4407
ab_1/cd_1	3261	$b_{r,2}$	4180
ab_2/cd_2	3180	$b_{r,3}$	4094
ab_3/cd_3	3076	$b_{r,4}$	3550
a_1/b_1	330	$a_{5,e}$	115
c_1	445	a_5	278
d_1	405	$b_{r,5,e}$	115
a_2/b_2	260	$b_{r,5}$	1420
c_2	400		
d_2	375		
a_3/b_3	325		
c_3	400		
d_3	370		
a_4/b_4	220		

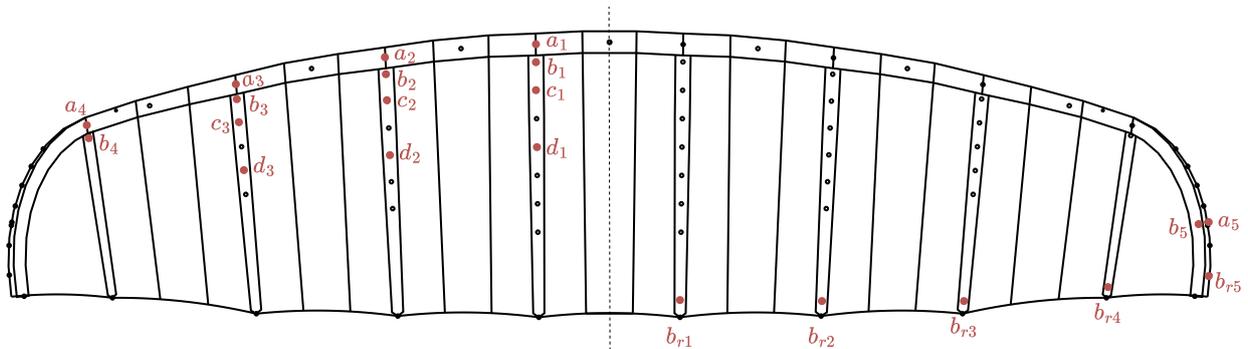


Figure F.1: Bridle attachment points on the wing. LE connections (left) correspond to the bridle system shown in Fig. F.2; TE connections (right) correspond to Fig. F.3.

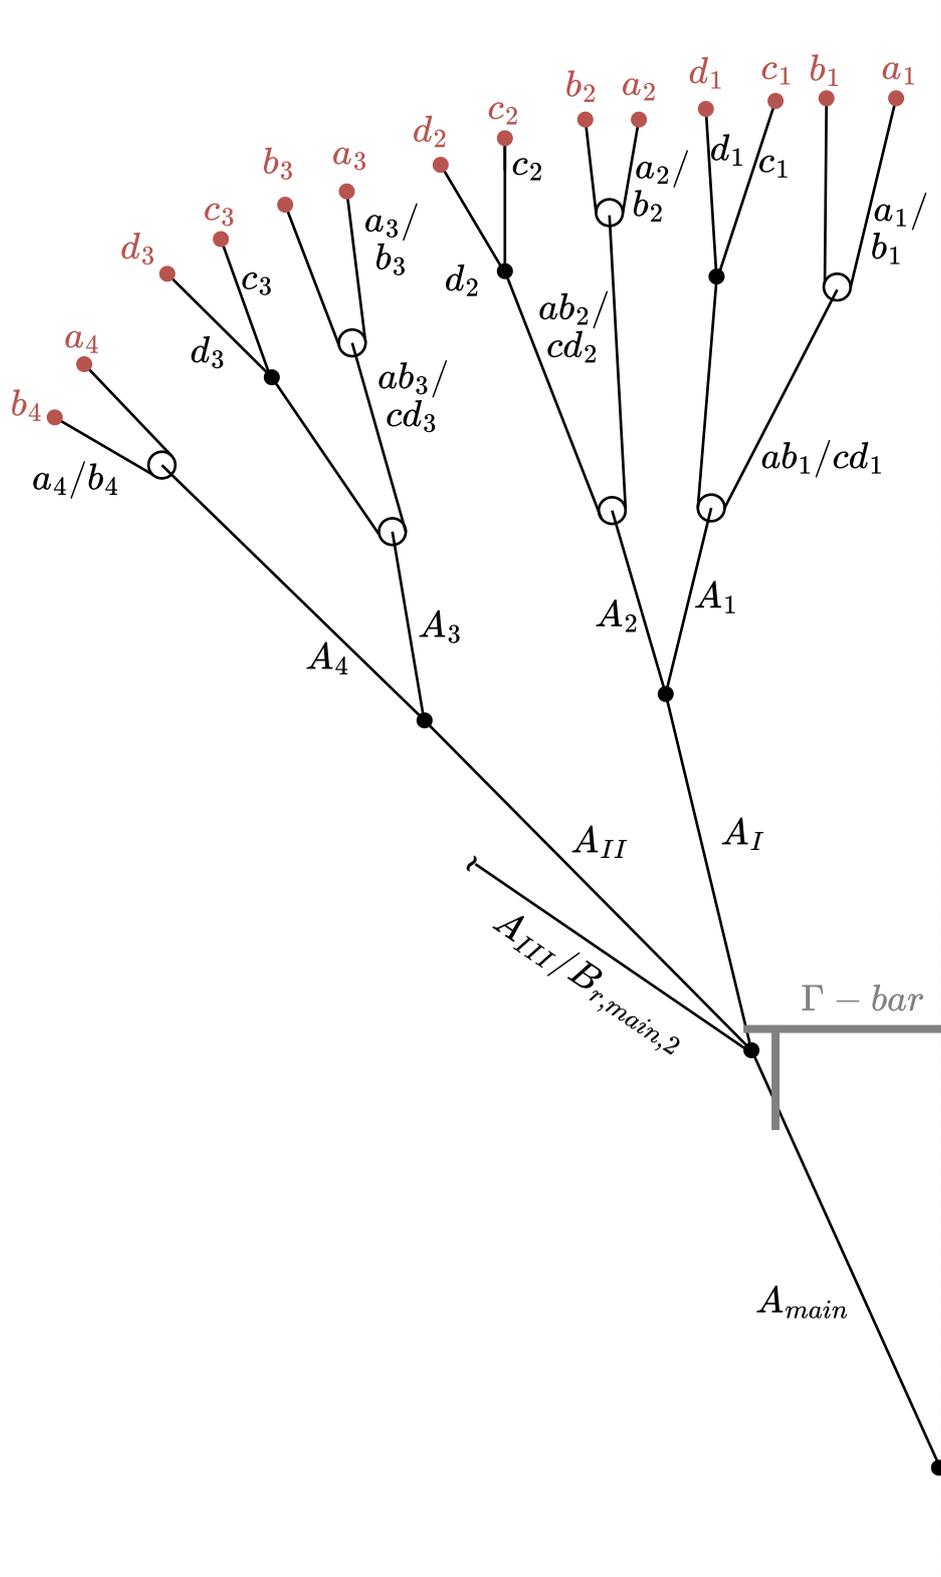


Figure F.2: Schematic of the leading-edge bridle system and corresponding segment numbering.

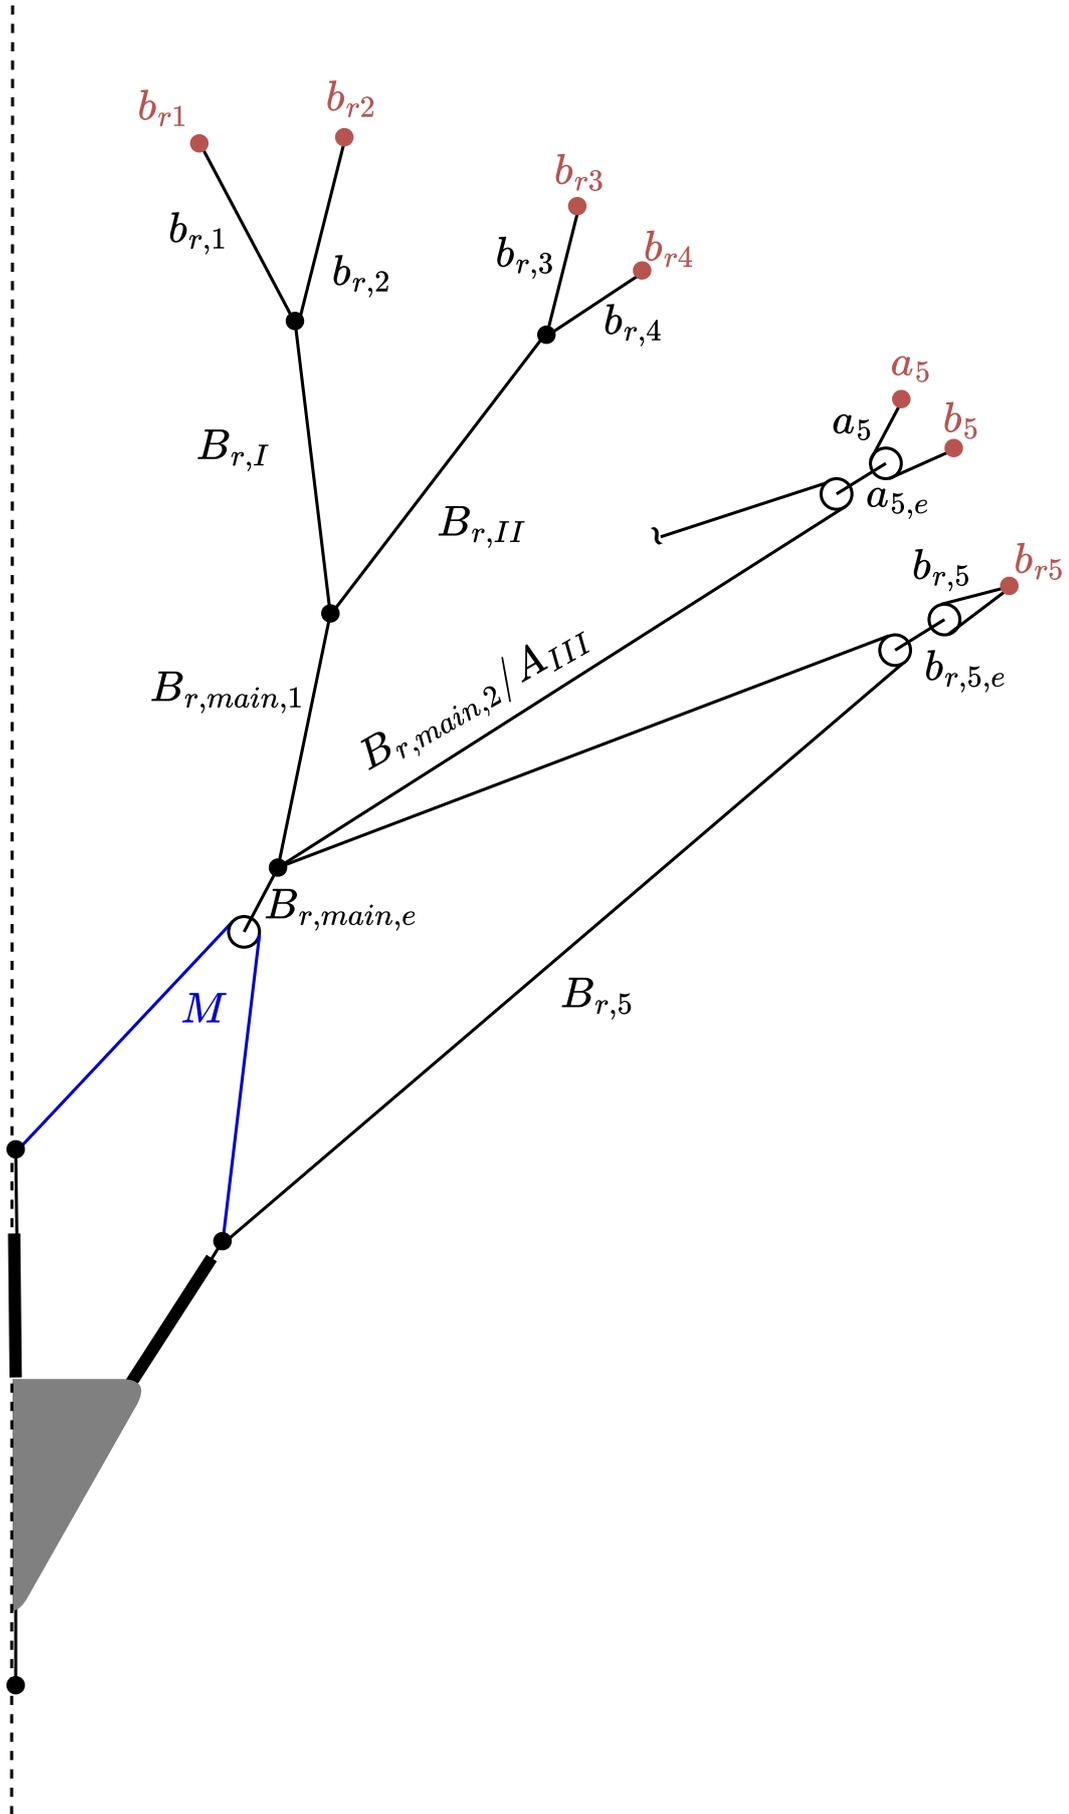


Figure F.3: Schematic of the trailing-edge bridle system and corresponding segment numbering.

Bridle geometry static test

This appendix presents the bridle configuration used during the static tests. For these tests, the bridle system was modified so that the kite could be suspended upside down from a ceiling with limited height. As a consequence, several bridle segments were shortened or omitted relative to the configuration used during flight testing. The rest of the bridle lengths are similar to the flight tests.

The resulting line lengths are listed in Table G.1. The numbering scheme for the front and rear bridle segments is shown in Fig. G.1 and Fig. G.2. The attachment points on the wing remain identical to the flight-test configuration, as illustrated in Fig. F.1. Black dots indicate knots here whilst white circles indicate pulleys.

Table G.1: Bridle length table used during the static experiments.

Line ID	Length (mm)	Line ID	Length (mm)
Front Lines		Rear Lines	
A_{main}	N/A	M	N/A
A_{I}	N/A	$B_{r,\text{main},e}$	70
A_{II}	3360	$B_{r,\text{main},1}$	610
$A_{\text{III}}/B_{r,\text{main},2}$	11420	$B_{r,\text{main},2}/A_{\text{III}}$	11420
$25A_1$	1860	$B_{r,I}$	2333
A_2	1798	$B_{r,II}$	2045
A_3	1791	$B_{r,5}$	9425
A_4	2877	$b_{r,1}$	4407
ab_1/cd_1	3261	$b_{r,2}$	4180
ab_2/cd_2	3180	$b_{r,3}$	4094
ab_3/cd_3	3076	$b_{r,4}$	3550
a_1/b_1	330	$a_{5,e}$	115
c_1	445	a_5	278
d_1	405	$b_{r,5,e}$	115
a_2/b_2	260	$b_{r,5}$	1420
c_2	400		
d_2	375		
a_3/b_3	325		
c_3	400		
d_3	370		
a_4/b_4	220		

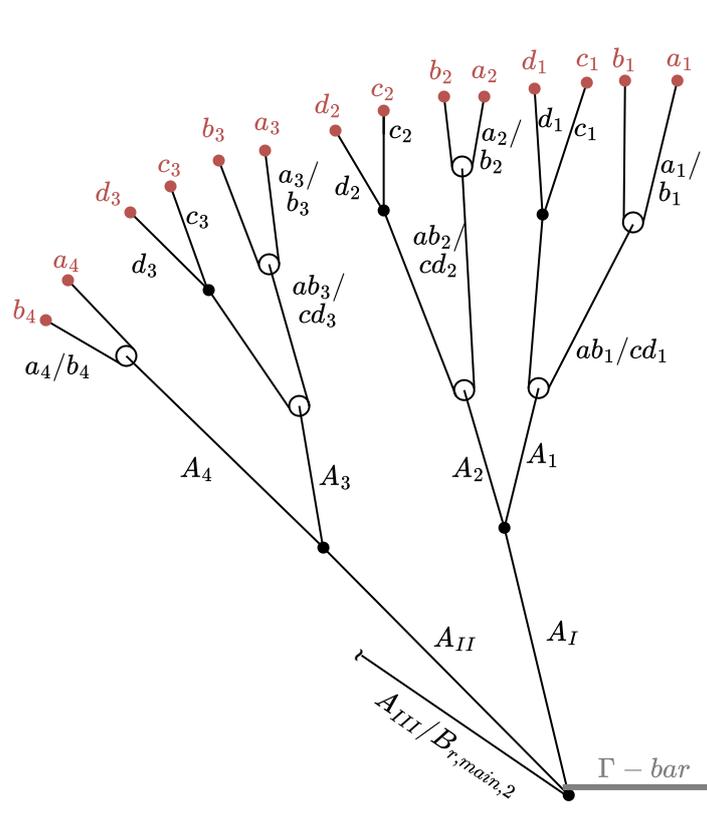


Figure G.1: Leading-edge bridle layout used during the static (upside-down) tests.

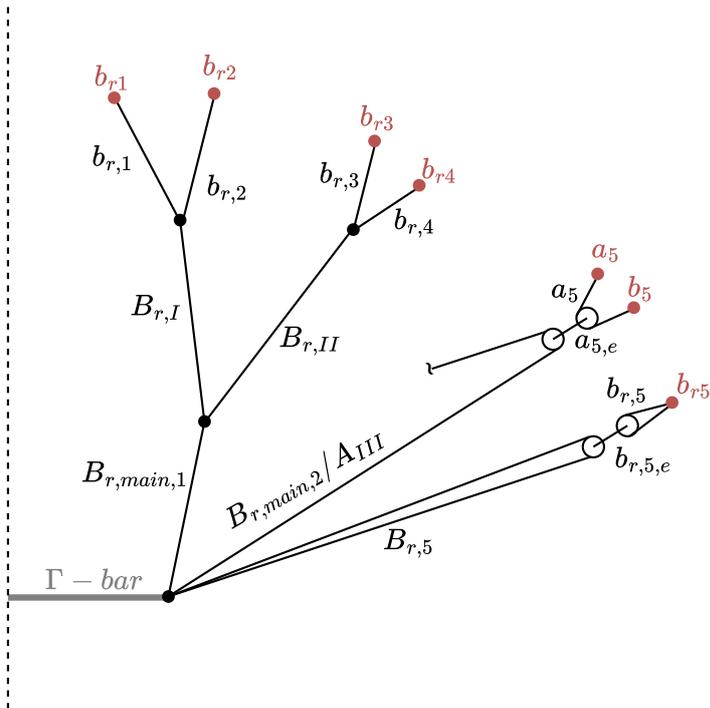


Figure G.2: Trailing-edge bridle layout used during the static (upside-down) tests.

Additional results H

This appendix presents some additional results complementary to results in the report. To this end, it shortly discusses a right turn in Appendix H.1, then shows a snapshot of the dynamic results in Appendix H.2. Lastly, the appendix ends with Appendix H.3.

H.1 Left turn

The shape-change results for a left-hand turn, complementing the right-turn case discussed in Chapter 9. The manoeuvre corresponds to a positive steering input u_s and produces a deformation pattern that is qualitatively consistent with the left turn, but mirrored in direction. Both bottom shear and front shear occur with reversed orientation. The spanwise contraction observed during the right turn is slightly smaller than in the left turn, which is consistent with the variability expected in real flight data, as well as the change in tether force.

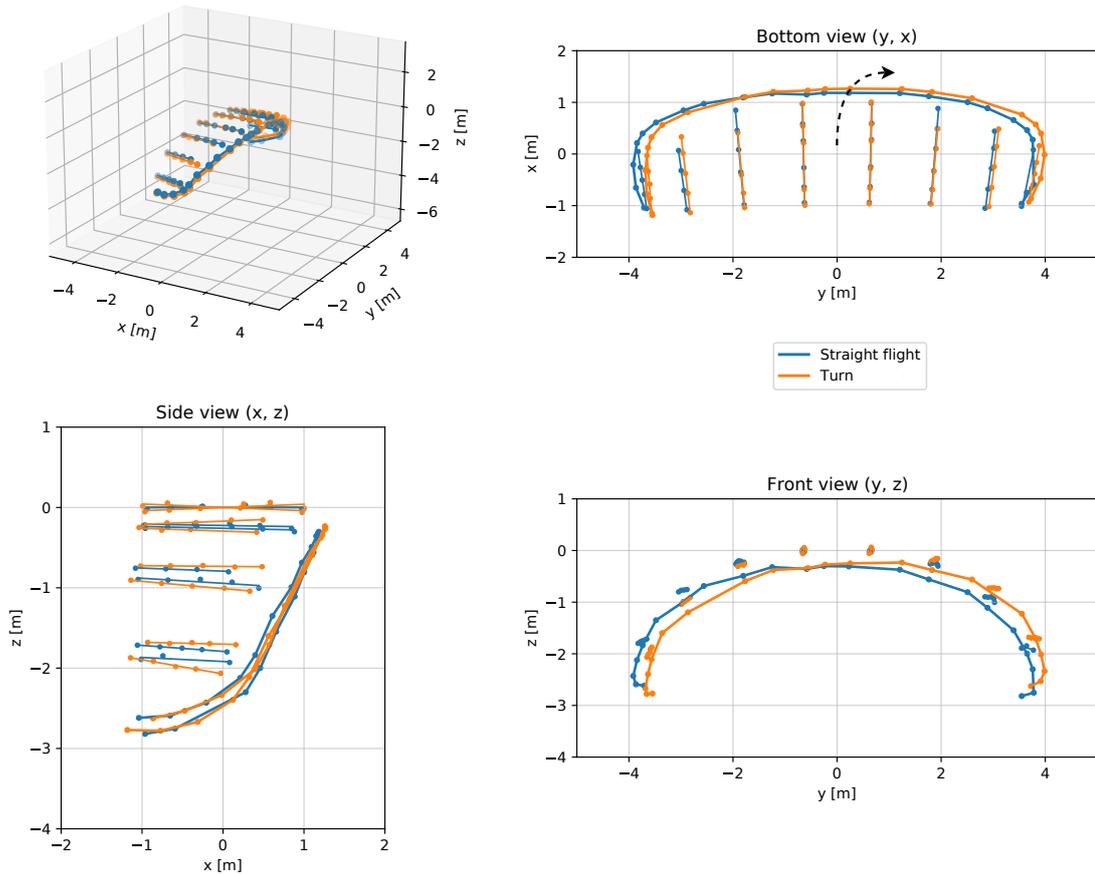


Figure H.1: Shape deformation during a left-hand turning manoeuvre.

Table H.1: Inputs and outputs corresponding to the two flight states shown in Fig. H.1.

State	u_{dp} (-)	u_s (-)	b_{UWB} (m)	V_a (ms^{-1})	α_{fl} ($^\circ$)	β_s ($^\circ$)	F_T/g (kg)	$\epsilon_{p,adj,avg}$ (cm)	$\epsilon_{p,adj,max}$ (cm)
Straight flight	0.3937	-0.0616	7.25	11.15	13.1	-4	110.97	5.58	8.21
Right turn	0.3937	0.1697	7.24	13.41	10.4	-5	283.73	3.38	5.23

H.2 Automated video results

Figure H.2 shows a representative frame from the automated post-processing video. The full video displays the reconstructed shape in time, along with four camera views, the synchronised flight state, and annotated geometric outputs.

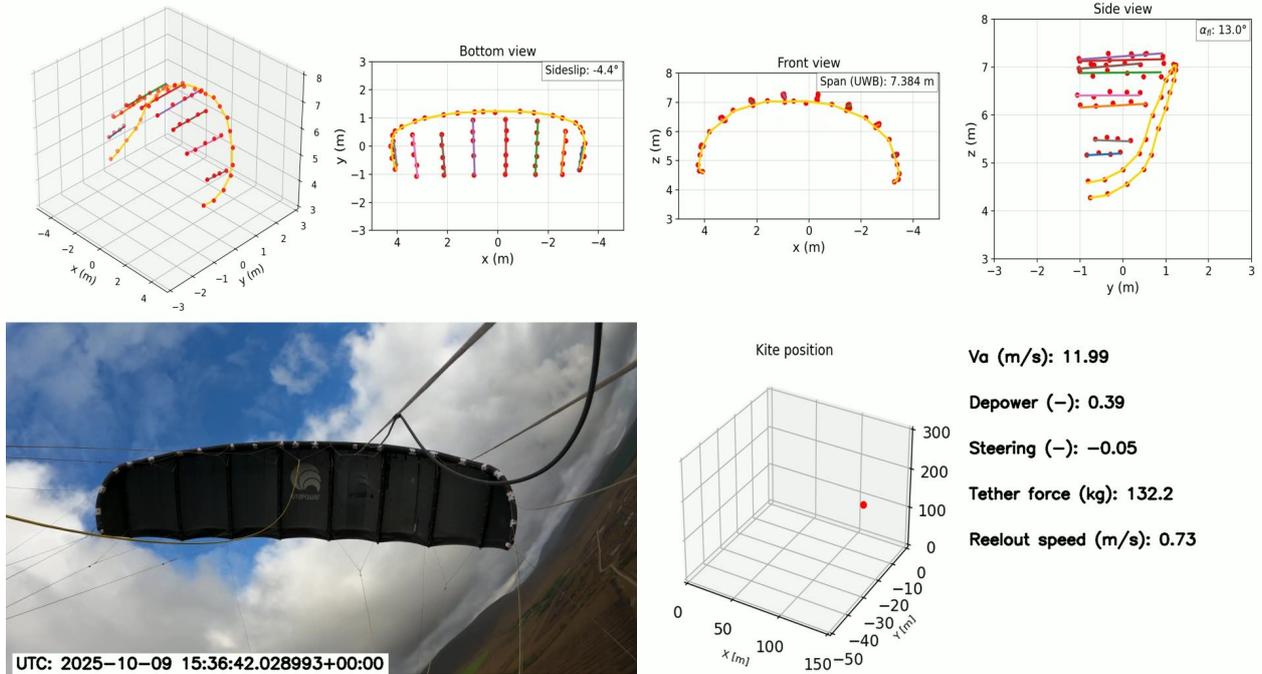


Figure H.2: Example snapshot from the automated result video, showing four point cloud views, synchronised video footage, and flight-state overlays.

H.3 Static test results

Figures H.3 and H.4 present the complete set of static test cases used for model validation, as discussed in Chapter 10. The results highlight the influence of bridle geometry, pressure level, and point-load configuration on the resulting wing shape.

In the unloaded configuration, the modified (shortened) bridle system used for suspension induces a noticeable inward deflection of the wing tips. Applying a small point load at the LE (PL1) produces only limited deformation. The larger LE point load (PL2) causes the centre struts to rotate upward significantly, resulting in a pitched appearance of the wing in the wing reference frame. Tip loads produce a geometry that more closely resembles the CAD reference, including outward-tilting wing tips.

While higher internal pressure reduces deformation, as expected, the magnitude of this effect is smaller than anticipated. This indicates that the overall stiffness of the system is governed predominantly by the bridle network rather than the internal pressure of the LE tube.

Kite shape for different load cases - Pressure 1

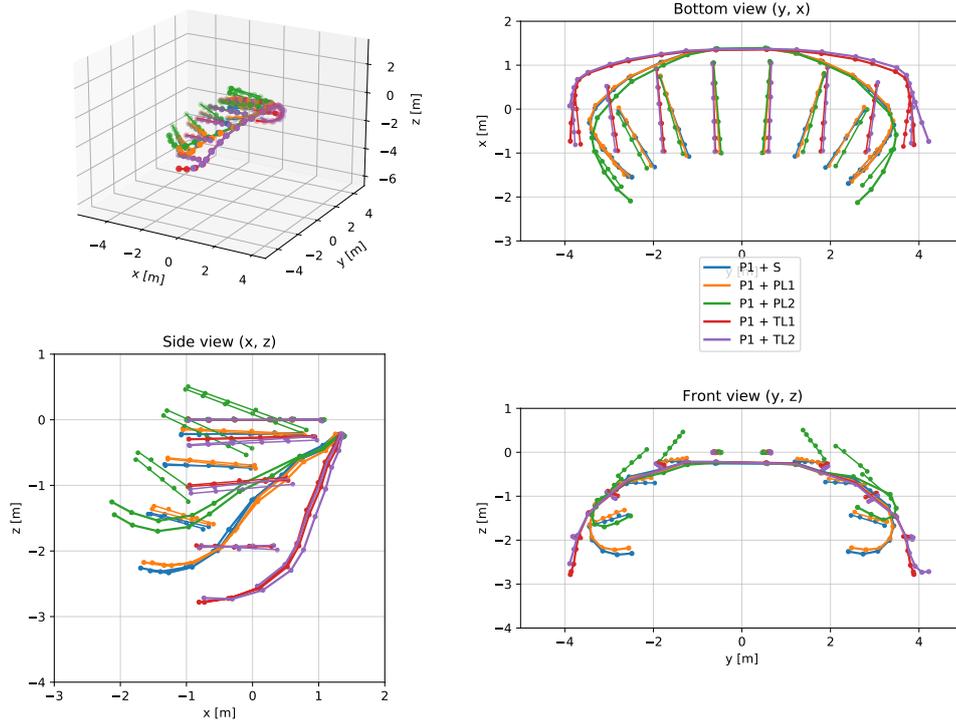


Figure H.3: Static test results for pressure level P1.

Kite shape for different load cases - Pressure 2

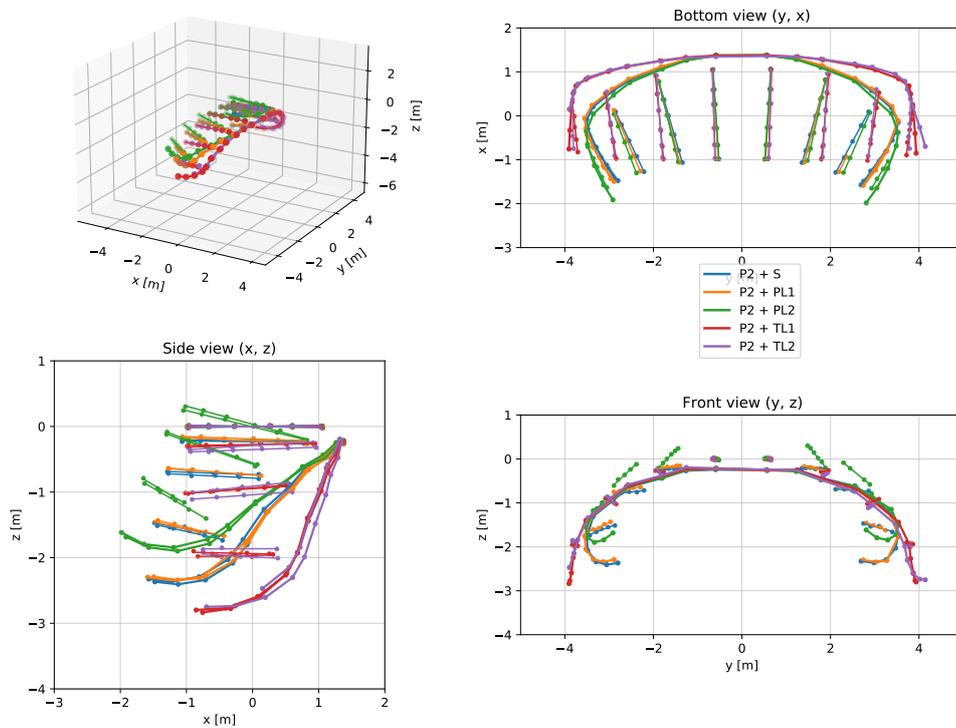


Figure H.4: Static test results for pressure level P2.

References

-
- [1] V. Masson-Delmotte *et al.*, “Global warming of 1.5 °c: An IPCC special report on the impacts of global warming of 1.5 °c above pre-industrial levels and related greenhouse gas emission pathways,” Intergovernmental Panel on Climate Change, Tech. Rep., 2018. <https://www.ipcc.ch/sr15/> (visited on 21/10/2025).
 - [2] J. Hansen *et al.*, “Ice melt, sea level rise and superstorms: Evidence from paleoclimate data, climate modeling, and modern observations that 2 °c global warming could be dangerous,” *Atmospheric Chemistry and Physics*, vol. 16, no. 6, pp. 3761–3812, 2016. doi: 10.5194/acp-16-3761-2016.
 - [3] K. Attanayake *et al.*, “Renewable energy as a solution to climate change: Insights from a comprehensive study across nations,” *PLoS One*, vol. 19, no. 6, e0299807, 2024. doi: 10.1371/journal.pone.0299807.
 - [4] L. Hagen, K. Petrick, S. Wilhelm, and R. Schmehl, “Life-cycle assessment of a multi-megawatt airborne wind energy system,” *Energies*, vol. 16, no. 4, 1750, 2023. doi: 10.3390/en16041750.
 - [5] Roland Schmehl, *Airborne wind energy: An introduction to an emerging technology*, 2019. <https://awesco.eu/awe-explained/> (visited on 29/3/2025).
 - [6] C. Grete, “Optimization, scaling and economics of pumping kite power systems,” MSc thesis, Delft University of Technology, The Netherlands, 2014. <https://resolver.tudelft.nl/uuid:b980aab7-f346-4030-97a3-a3cde13a51d6>.
 - [7] M. L. Loyd, “Crosswind kite power,” *Journal of Energy*, vol. 4, no. 3, pp. 106–111, 1980. doi: 10.2514/3.48021.
 - [8] R. van der Vlugt, J. Peschel, and R. Schmehl, “Design and experimental characterization of a pumping kite power system,” in *Airborne Wind Energy*, U. Ahrens, M. Diehl, and R. Schmehl, Eds., Springer, 2013, pp. 403–425. doi: 10.1007/978-3-642-39965-7_23.
 - [9] T. N. Dief, U. Fechner, R. Schmehl, S. Yoshida, and M. A. Rushdi, “Adaptive flight path control of airborne wind energy systems,” *Energies*, vol. 13, no. 3, 667, 2020. doi: 10.3390/en13030667.
 - [10] J. A. W. Poland and R. Schmehl, “Modelling aero-structural deformation of flexible membrane kites,” *Energies*, vol. 16, no. 14, 5264, 2023. doi: 10.3390/en16145264.
 - [11] Kitepower B.V., *Kitepower airborne wind energy-plug and play, mobile wind energy*, 2025. <https://thekitepower.com/> (visited on 26/8/2025).
 - [12] J. Degroote, K.-J. Bathe, and J. Vierendeels, “Performance of a new partitioned procedure versus a monolithic procedure in fluid–structure interaction,” *Computers and Structures*, vol. 87, no. 11–12, pp. 793–801, 2009. doi: 10.1016/j.compstruc.2008.11.013.
 - [13] O. Cayon, M. Gaunaa, and R. Schmehl, “Fast aero-structural model of a leading-edge inflatable kite,” *Energies*, vol. 16, no. 7, 3061, 2023. doi: 10.3390/en16073061.
 - [14] J. A. W. Poland and R. Schmehl, “A virtual wind tunnel for deforming airborne wind energy kites,” *Journal of Physics: Conference Series*, vol. 2767, no. 7, 072001, 2024. doi: 10.1088/1742-6596/2767/7/072001.
 - [15] S. Jonard, “In-flight measurement of kite deformations with inertial sensors,” MSc thesis, Delft University of Technology, The Netherlands, 2021. <https://resolver.tudelft.nl/uuid:c2db7fb0-ddb9-4bc0-af41-e5cca64d4f58>.
 - [16] J. Oehler and R. Schmehl, “Aerodynamic characterization of a soft kite by in situ flow measurement,” *Wind Energy Science*, vol. 4, no. 1, pp. 1–21, 2019. doi: 10.5194/wes-4-1-2019.
 - [17] J. Breukels, “An engineering methodology for kite design,” PhD thesis, Delft University of Technology, The Netherlands, 2011. <https://resolver.tudelft.nl/uuid:cdece38a-1f13-47cc-b277-ed64fdda7cdf>.
 - [18] N. Landel-Mills, *Kitesurfing explained by newtonian physics: The kite steals momentum from the wind*, 2021. doi: 10.13140/RG.2.2.28835.37929.
 - [19] S. Becker and P. Bruce, “Experimental study of paraglider aerodynamics,” MSc thesis, Imperial College London, United Kingdom, 2017. doi: 10.13140/RG.2.2.33674.16321.

- [20] N. A. Fogell, S. J. Sherwin, C. J. Cotter, L. Iannucci, R. Palacios, and D. Pope, “Fluid-structure interaction simulation of the inflated shape of ram-air parachutes,” in *AIAA Aerodynamic Decelerator Systems (ADS) Conference 2013*, 2013, 1326. doi: 10.2514/6.2013-1326.
- [21] M. Folkersma, R. Schmehl, and A. Viré, “Boundary layer transition modeling on leading edge inflatable kite airfoils,” *Wind Energy*, vol. 22, no. 7, pp. 908–921, 2019. doi: 10.1002/we.2329.
- [22] J. Watchorn, “Aerodynamic load modelling for leading edge inflatable kites,” MSc thesis, Delft University of Technology, The Netherlands, 2023. <https://resolver.tudelft.nl/uuid:42f611a2-ef79-4540-a43c-0ea827700388>.
- [23] P. Demkowicz, “Numerical analysis of the flow past a leading edge inflatable kite wing using a correlation-based transition model,” MSc thesis, Delft University of Technology, The Netherlands, 2019. <https://resolver.tudelft.nl/uuid:c53aa605-1b2e-47a7-b991-c1917d7463b4>.
- [24] M. Folkersma, “Aeroelasticity of membrane kites: Airborne wind energy applications,” PhD thesis, Delft University of Technology, The Netherlands, 2022. doi: 10.4233/uuid:cae39f5a-49bc-438b-948f-b6ab51208068.
- [25] G. H. M. Lebesque, “Steady-state rans simulation of a leading edge inflatable wing with chordwise struts,” MSc thesis, Delft University of Technology, The Netherlands, 2020. <https://resolver.tudelft.nl/uuid:f0bc8a1e-088d-49c5-9b77-ebf9e31cf58b>.
- [26] P. Rojratsirikul, “Aerodynamics of flexible membranes,” PhD thesis, University of Bath, United Kingdom, 2010. <https://researchportal.bath.ac.uk/en/studentTheses/aerodynamics-of-flexible-membranes/>.
- [27] A. de Wachter, “Deformation and aerodynamic performance of a ram-air wing,” MSc thesis, Delft University of Technology, The Netherlands, 2008. <https://resolver.tudelft.nl/uuid:786e3395-4590-4755-829f-51283a8df3d2>.
- [28] T. van Lith, “Aerodynamic analysis of a 2d rigid lei airfoil,” MSc thesis, Delft University of Technology, The Netherlands, 2025. <https://resolver.tudelft.nl/uuid:53c595f1-2beb-4cc3-853d-f3a51a116690>.
- [29] R. Schmehl, “Airborne wind energy: Physics of tethered flight,” Lecture slides, TU Delft, Oct. 2023.
- [30] U. Fechner and R. Schmehl, “Flight path planning in a turbulent wind environment,” in *Airborne Wind Energy*, R. Schmehl, Ed., Springer, 2018, pp. 361–390. doi: 10.1007/978-981-10-1947-0_15.
- [31] R. C. Leuthold, “Multiple-wake vortex lattice method for airborne wind energy membrane-wing kites,” MSc thesis: a literature review, Delft University of Technology, The Netherlands, 2015. <http://resolver.tudelft.nl/uuid:bad76ade-81d3-4d86-a631-c45319812445>.
- [32] K. Alexander and J. Stevenson, “Kite equilibrium and bridle length,” *The Aeronautical Journal*, vol. 105, no. 1051, pp. 535–541, 2001. doi: 10.1017/S0001924000017991.
- [33] M. Schelbergen and R. Schmehl, “Swinging motion of a kite with suspended control unit flying turning manoeuvres,” *Wind Energy Science*, vol. 9, no. 6, pp. 1323–1344, 2024. doi: 10.5194/wes-9-1323-2024.
- [34] J. A. W. Poland, “Modelling aeroelastic deformation of soft wing membrane kites,” MSc thesis, Delft University of Technology, The Netherlands, 2022. <https://resolver.tudelft.nl/uuid:39d67249-53c9-47b4-84c0-ddac948413a5>.
- [35] Deutsche Welle. “Sky-high kites aim to tap unused wind power.” (2023), <https://p.dw.com/p/4Op36> (visited on 27/5/2025).
- [36] H. A. Bosch, “Finite element analysis of a kite for power generation,” MSc thesis, Delft University of Technology, The Netherlands, 2012. <https://resolver.tudelft.nl/uuid:888fe64a-b101-438c-aa6f-8a0b34603f8e>.
- [37] M. Erhard and H. Strauch, “Theory and experimental validation of a simple comprehensible model of tethered kite dynamics used for controller design,” in *Airborne Wind Energy*, U. Ahrens, M. Diehl, and R. Schmehl, Eds., Springer, 2022, pp. 141–165. doi: 10.1007/978-3-642-39965-7_8.
- [38] C. Elfert, D. Göhlich, and R. Schmehl, “Measurement of the turning behaviour of tethered membrane wings using automated flight manoeuvres,” *Wind Energy Science*, vol. 9, no. 12, pp. 2261–2282, 2024. doi: 10.5194/wes-9-2261-2024.
- [39] R. Leuthold, “Multiple-wake vortex lattice method for membrane-wing kites,” MSc thesis, Delft University of Technology, The Netherlands, 2015. <https://resolver.tudelft.nl/uuid:4c2f34c2-d465-491a-aa64-d991978fedf4>.
- [40] H. Nyquist, “Certain topics in telegraph transmission theory,” *Transactions of the American Institute of Electrical Engineers*, vol. 47, no. 2, pp. 617–644, 1928. doi: 10.1109/T-AIEE.1928.5055024.

- [41] C. E. Shannon, “Communication in the presence of noise,” *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, 1949. doi: 10.1109/JRPROC.1949.232969.
- [42] B. Perin, A. Donnard, P. Bordenave, C. Larrieu, C. Simond, and H. Belloc, “Fluid-structure interaction simulation of ram air parachutes - an application for a kite,” in *Aerodynamic Decelerator Systems Technology Conferences*, 2015. doi: 10.2514/6.2015-2185.
- [43] J. F. J. E. M. Schwoel, “Finite element analysis of inflatable structures using uniform pressure: Inflatable beam validation and leading edge inflatable tube kite structural modeling in madymo,” MSc thesis, Delft University of Technology, The Netherlands, 2012. <https://resolver.tudelft.nl/uuid:f92da57f-55df-4109-9f8a-8c7c2b220c6a>.
- [44] K. Hanke and S. Schenk, “Evaluating the geometric shape of a flying paraglider,” *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XL-5, pp. 265–272, 2014. doi: 10.5194/isprsarchives-XL-5-265-2014.
- [45] M. Behrel, “Les kites comme propulsion auxiliaire pour les navires: Système expérimental, campagnes de mesures, analyse des données et identification des performances des kites,” PhD thesis, Université de Bretagne occidentale, France, 2017. <https://theses.hal.science/tel-03081135>.
- [46] R. Infanzon, “Photogrammetry based investigation of a kite’s flying shape,” National Centre for Maritime Engineering and Hydrodynamics, Australian Maritime College, University of Tasmania, Launceston, TAS, Australia, Technical Report, 2013, Not publicly available.
- [47] D. Costa, “Experimental investigation of aerodynamic and structural properties of a kite,” Not publicly available, MSc thesis, Eidgenössische Technische Hochschule Zürich, Switzerland, 2011.
- [48] M. R. van Reijen, “The turning of kites: A quantification of known theories,” MSc thesis, Delft University of Technology, The Netherlands, 2018. <https://resolver.tudelft.nl/uuid:5836c754-68d3-477a-be32-8e1878f85eac>.
- [49] J. Oehler, “Measuring apparent flow vector on a flexible wing kite,” MSc thesis, Universität Stuttgart, Germany, 2017. <http://elib.uni-stuttgart.de/handle/11682/9907>.
- [50] R. Borobia-Moreno, D. Ramiro-Rebollo, R. Schmehl, and G. Sánchez-Arriaga, “Identification of kite aerodynamic characteristics using the estimation before modeling technique,” *Wind Energy*, vol. 24, no. 6, pp. 596–608, 2021. doi: 10.1002/we.2591.
- [51] D. J. Thevara and C. Vasanth Kumar, “Application of photogrammetry to automated finishing operations,” *IOP Conference Series: Materials Science and Engineering*, vol. 402, no. 1, 012025, 2018. doi: 10.1088/1757-899X/402/1/012025.
- [52] K. Murawski, “Method of measuring the distance to an object based on one shot obtained from a motionless camera with a fixed-focus lens,” *Acta Physica Polonica A*, vol. 127, no. 6, pp. 1591–1596, 2015. doi: 10.12693/APhysPolA.127.1591.
- [53] K. Yue, Z. Li, M. Zhang, and S. Chen, “Transient full-field vibration measurement using spectroscopical stereo photogrammetry,” *Optics Express*, vol. 18, no. 26, pp. 26 866–2687, 2010. doi: 10.1364/OE.18.026866.
- [54] Z. Kalal, K. Mikolajczyk, and J. Matas, “Tracking-learning-detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 7, pp. 1409–1422, 2012. doi: 10.1109/TPAMI.2011.239.
- [55] M. Chapel and T. Bouwmans, “Moving objects detection with a moving camera: A comprehensive review,” *Computer Science Review*, vol. 38, 100310, 2020. doi: 10.1016/j.cosrev.2020.100310.
- [56] C. Zheng, Y. Ge, and A. Guo, *Ultra-wideband technology: Characteristics, applications and challenges*, 2023. doi: 10.48550/arXiv.2307.13066. arXiv: 2307.13066 [eess.SP].
- [57] M. Cheraghinia *et al.*, “A comprehensive overview on uwb radar: Applications, standards, signal processing techniques, datasets, radio chips, trends and future research directions,” *IEEE Communications Surveys & Tutorials*, vol. 27, no. 4, pp. 2283–2324, 2025. doi: 10.1109/COMST.2024.3488173.
- [58] A. Alarifi *et al.*, “Ultra wideband indoor positioning technologies: Analysis and recent advances,” *Sensors*, vol. 16, no. 5, 707, 2016. doi: 10.3390/s16050707.
- [59] B. Großwindhager *et al.*, “Salma: Uwb-based single-anchor localization system using multipath assistance,” in *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems (SenSys’18)*, 2018, pp. 132–144. doi: 10.1145/3274783.3274844.
- [60] T. Polonelli, S. Schläpfer, and M. Magno, “Performance comparison between decawave dw1000 and dw3000 in low-power double side ranging applications,” in *2022 IEEE Sensors Applications Symposium (SAS)*, 2022, pp. 1–6. doi: 10.1109/SAS54819.2022.9881375.

- [61] Y. Liu and Y. Bao, “Real-time remote measurement of distance using ultra-wideband (uwb) sensors,” *Automation in Construction*, vol. 150, 104849, 2023. doi: 10.1016/j.autcon.2023.104849.
- [62] G. Burgess, M. Shortis, and P. Scott, “Photographic assessment of retroreflective film properties,” *Isprs Journal of Photogrammetry and Remote Sensing*, vol. 66, no. 5, pp. 743–750, 2011. doi: 10.1016/j.isprsjprs.2011.07.002.
- [63] M. Anusha, C. Vidya Raj, and H. P. Vinod Kumar, “Integration of all types of photoelectric sensors with iot and development of android based application,” *International Research Journal of Engineering and Technology (IRJET)*, vol. 8, no. 8, pp. 1603–1608, 2021. <https://www.irjet.net/archives/V8/i8/IRJET-V8I8229.pdf> (visited on 1/4/2025).
- [64] B. Fei, J. Xu, R. Zhang, Q. Zhou, W. Yang, and Y. He, “3d gaussian splatting as new era: A survey,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 8, pp. 1–20, 2024. doi: 10.1109/TVCG.2024.3397828.
- [65] C. Jiang, R. Gao, K. Shao, Y. Wang, R. Xiong, and Y. Zhang, *Li-gs: Gaussian splatting with lidar incorporated for accurate large-scale reconstruction*, 2024. doi: 10.48550/arXiv.2409.12899. arXiv: 2409.12899 [cs.RG].
- [66] J. Raychev, G. Hristov, D. Kyuchukova, and P. Zahariev, “Workflow for development of a virtual museum that will provide better way for learning the cultural heritage,” in *16th International Conference on Information Technology Based Higher Education and Training (ITHET)*, 2017, pp. 1–5. doi: 10.1109/ITHET.2017.8067815.
- [67] A. Georgopoulos, C. Ioannidis, and A. Valanis, “Assessing the performance of a structured light scanner,” *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 38, no. 5, pp. 250–255, 2010. <https://www.isprs.org/proceedings/xxxviii/part5/papers/177.pdf> (visited on 23/4/2025).
- [68] J. Schipper *et al.*, “Reliability and validity of handheld structured light scanners and a static stereophotogrammetry system in facial three-dimensional surface imaging,” *Scientific Reports*, vol. 14, 8172, 2024. doi: 10.1038/s41598-024-57370-x.
- [69] A. Ignatov *et al.*, *Fast and accurate single-image depth estimation on mobile devices, mobile ai 2021 challenge: Report*, 2021. doi: 10.48550/arXiv.2105.08630. arXiv: 2105.08630 [eess.IV].
- [70] V. Kumar Jha *et al.*, “A depth image quality benchmark of three popular low-cost depth cameras,” *MM Science Journal*, vol. 2020, no. 5, pp. 4194–4200, 2020. doi: 10.17973/MMSJ.2020_12_2020057.
- [71] L. Li, “Time-of-flight camera – an introduction,” Texas Instruments, Tech. Rep. SLOA190B, 2014. <https://www.ti.com/lit/wp/sloa190b/sloa190b.pdf?ts=1765358562047> (visited on 19/3/2025).
- [72] ToF Sensors, *Comparison between tof sensors and other types of distance sensors*, 2023. <https://tofensors.com/comparison-between-tof-sensors-and-other-types-of-distance-sensors> (visited on 29/3/2025).
- [73] M. Hansard, S. Lee, O. Choi, and R. Horaud, *Time-of-Flight Cameras*, 1st ed. Springer, 2012, ISBN: 978-1-4471-4658-2. doi: 10.1007/978-1-4471-4658-2.
- [74] D. A. Thornton, K. Redmill, and B. Coifman, “Automated parking surveys from a lidar equipped vehicle,” *Transportation Research Part C: Emerging Technologies*, vol. 39, no. 3, pp. 23–35, 2014. doi: 10.1016/j.trc.2013.11.014.
- [75] C. Benedek, A. Majdik, B. Nagy, Z. Rozsa, and T. Sziranyi, “Positioning and perception in lidar point clouds,” *Digital Signal Processing*, vol. 119, no. 3, 103193, 2021. doi: 10.1234/abcd.efgh.
- [76] Y. Waykar, “Lidar technology: A comprehensive review and future prospects,” *Journal of Emerging Technologies and Innovative Research*, vol. 9, no. 7, pp. 164–172, 2022. <http://www.jetir.org/papers/JETIR2207728.pdf> (visited on 20/3/2025).
- [77] B. Schlager, T. Goelles, M. Behmer, S. Muckenhuber, J. Payer, and D. Watzenig, “Automotive lidar and vibration: Resonance, inertial measurement unit, and effects on the point cloud,” *IEEE Open Journal of Intelligent Transportation Systems*, vol. 3, pp. 426–434, 2022. doi: 10.1109/OJITS.2022.3176471.
- [78] H. W. Yoo, R. Riegler, D. Brunner, S. Albert, T. Thurner, and G. Schitter, “Experimental evaluation of vibration influence on a resonant mems scanning system for automotive lidars,” *IEEE Transactions on Industrial Electronics*, vol. 69, no. 3, pp. 3099–3108, 2022. doi: 10.1109/tie.2021.3065608.
- [79] J. C. Fernandez Diaz, W. E. Carter, R. L. Shrestha, and C. L. Glennie, “Lidar remote sensing,” in *Handbook of Satellite Applications*, J. N. Pelton, S. Madry, and S. Camacho-Lara, Eds., Springer International Publishing, 2017, pp. 929–980. doi: 10.1007/978-3-319-23386-4_44.

- [80] G. P. Yolanda, T. Guzmán, and T. González, “Development of a low-cost, short-range radar system to measure speed and distance,” *TECCIENCIA*, vol. 12, no. 22, pp. 99–106, 2017. doi: 10.18180/teccien-
cia.2017.22.11.
- [81] N. Levanon and I. Cohen, “Bistatic radar positioning of a moving target, using only range and range-rate measurements,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 59, no. 2, pp. 1581–1589, 2023. doi: 10.1109/TAES.2022.3202758.
- [82] O. Kanhere, S. Goyal, M. Beluri, and T. S. Rappaport, “Target localization using bistatic and multistatic radar with 5g nr waveform,” in *2021 IEEE 93rd Vehicular Technology Conference (VTC-Spring)*, 2021, pp. 1–7. doi: 10.1109/VTC2021-Spring51267.2021.9449071.
- [83] Analog Devices, Inc., *Select the best adc for radar phased-array apps — part 1*, 2023. <https://www.analog.com/en/resources/technical-articles/select-the-best-adc-for-radar-phased-array-apps-part1.html> (visited on 3/4/2025).
- [84] M. Toa and A. Whitehead, “Ultrasonic sensing basics,” Texas Instruments, Tech. Rep. SLAA907D, 2020, pp. 53–75. <https://www.ti.com/lit/an/slaa907d/slaa907d.pdf> (visited on 28/3/2025).
- [85] L. Koval, J. Vaňuš, and P. Bílík, “Distance measuring by ultrasonic sensor,” *IFAC-PapersOnLine*, vol. 49, no. 25, pp. 153–158, 2016. doi: 10.1016/j.ifacol.2016.12.026.
- [86] E. Sze, D. Hindarto, I. K. A. Wirayasa, and Haryono, “Performance comparison of ultrasonic sensor accuracy in measuring distance,” *Sinkron: Jurnal dan Penelitian Teknik Informatika*, vol. 7, no. 4, pp. 2556–2562, 2022. doi: 10.33395/sinkron.v7i4.11883.
- [87] Z. Qiu, Y. Lu, and Z. Qiu, “Review of ultrasonic ranging methods and their current challenges,” *Micromachines*, vol. 13, no. 4, 520, 2022. doi: 10.3390/mi13040520.
- [88] M. Papoutsidakis, C. Psomopoulos, G. Ioannidis, and D. Tseles, “Motion sensors and transducers to navigate an intelligent mechatronic platform for outdoor applications,” *Sensors and Transducers*, vol. 198, no. 3, pp. 16–24, 2016. https://www.researchgate.net/publication/299511206_Motion_Sensors_and_Transducers_to_Navigate_an_Intelligent_Mechatronic_Platform_for_Outdoor_Applications (visited on 9/3/2025).
- [89] A. Shrivastava, A. Verma, and S. Singh, “Effect of variation of separation between the ultrasonic transmitter and receiver on the accuracy of distance measurement,” *International Journal of Computer Science and Information Technology*, vol. 1, no. 2, pp. 19–28, 2009. <https://airccse.org/journal/ijcsitcurr.html> (visited on 10/3/2025).
- [90] M. R. Kaloop, E. Elbeltagi, J. W. Hu, and A. Elrefai, “Recent advances of structures monitoring and evaluation using gps-time series monitoring systems: A review,” *ISPRS International Journal of Geo-Information*, vol. 6, no. 12, 382, 2017. doi: 10.3390/ijgi6120382.
- [91] N. M. Drawil, H. M. Amar, and O. A. Basir, “Gps localization accuracy classification: A context-based approach,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 1, pp. 262–273, 2013. doi: 10.1109/TITS.2012.2213815.
- [92] C. Mekik and M. Arslanoglu, “Investigation on accuracies of real time kinematic gps for gis applications,” *Remote Sensing*, vol. 1, no. 1, pp. 22–35, 2009. doi: 10.3390/rs1010022.
- [93] D. Tran and L. Dũng, “Study on the positioning efficiency of gnss rtk for road profile surveys - case study in vietnam,” *Journal of Science and Technology in Civil Engineering (JSTCE) - HUCE*, vol. 18, no. 2, pp. 86–98, 2024. doi: 10.31814/stce.huce2024-18(2)-07.
- [94] G. W. Roberts, X. Tang, and X. He, “Accuracy analysis of gps/bds relative positioning using zero-baseline measurements,” *The Journal of Global Positioning Systems*, vol. 16, no. 7, 7, 2018. doi: 10.1186/s41445-018-0015-6.
- [95] G. Frigyes, E. Myers, and J. Allison, *Fundamentals of photoelectric sensors*, 2010. <https://www.automation.com/en-us/articles/2014-1/fundamentals-of-photoelectric-sensors> (visited on 3/6/2025).
- [96] J. Hulth, “Using a draw-wire sensor to continuously monitor glacier melt,” *Journal of Glaciology*, vol. 56, no. 199, pp. 922–924, 2010. doi: 10.3189/002214310794457290.
- [97] J. Song, P. Si, H. Hua, and M. Qiu, “Research on the inherent nonlinearity calibration of the potentiometer of a miniature linear series elastic actuator,” *Actuators*, vol. 11, no. 8, 207, 2022. doi: 10.3390/act11080207.
- [98] N. Bin Ahmad, R. A. Raja Ghazilla, N. Khairi, and V. Kasi, “Reviews on various inertial measurement unit (imu) sensor applications,” *International Journal of Signal Processing Systems*, vol. 1, no. 2, pp. 256–262, 2013. doi: 10.12720/ijsp.1.2.256-262.

- [99] P. Franček, K. Jambrošić, M. Horvat, and V. Planinec, “The performance of inertial measurement unit sensors on various hardware platforms for binaural head-tracking applications,” *Sensors*, vol. 23, no. 2, 872, 2023. doi: 10.3390/s23020872.
- [100] Xsens, *Xsens dot user manual*, 2025. <https://www.xsens.com/hubfs/Downloads/Manuals/Xsens%20DOT%20User%20Manual.pdf> (visited on 26/2/2025).
- [101] J. Neilson, P. Cataldi, and B. Derby, *Graphene-based transparent flexible strain gauges with tunable sensitivity and strain range*, 2023. arXiv: 2304.14297 [physics.app-ph].
- [102] M. Elsherif *et al.*, “Optical fiber sensors: Working principle, applications, and limitations,” *Advanced Photonics Research*, vol. 3, no. 11, 2100371, 2022. doi: 10.1002/adpr.202100371.
- [103] M. Ju *et al.*, “Piezoelectric materials and sensors for structural health monitoring: Fundamental aspects, current status, and future perspectives,” *Sensors*, vol. 23, no. 1, 543, 2023. doi: 10.3390/s23010543.
- [104] D. A. Litteken, “Evaluation of strain measurement devices for inflatable structures,” in *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference (AIAA SciTech 2017)*, 2017. doi: 10.2514/6.2017-0426.
- [105] H. Nesser, H. A. Mahmoud, and G. Lubineau, “High-sensitivity rfid sensor for structural health monitoring,” *Advanced Science*, vol. 10, no. 26, 2301807, 2023. doi: 10.1002/advs.202301807.
- [106] L. Bertram, M. Brink, and W. Lang, “Wireless, material-integrated sensors for strain and temperature measurement in glass fibre reinforced composites,” *Sensors*, vol. 23, no. 14, 6375, 2023. doi: 10.3390/s23146375.
- [107] B. Wang, X. Fan, Y. Fu, and Z. He, “Dynamic strain measurement with khz-level repetition rate and centimeter-level spatial resolution based on brillouin optical correlation domain analysis,” *Optics Express*, vol. 26, no. 6, pp. 6916–6928, 2018. doi: 10.1364/OE.26.006916.
- [108] P. Yuxin *et al.*, “Recent advances in flexible bending sensors and their applications,” *International Journal of Smart and Nano Materials*, vol. 15, no. 4, pp. 697–729, 2024. doi: 10.1080/19475411.2024.2417834.
- [109] B. Zazoum, K. M. Battoo, and M. A. A. Khan, “Recent advances in flexible sensors and their applications,” *Sensors*, vol. 22, no. 12, 4653, 2022. doi: 10.3390/s22124653.
- [110] C. Zhang, S. Jiang, H. Wang, J. Liu, A. Jain, and M. Armand, *Data-driven shape sensing in continuum manipulators via sliding resistive flex sensors*, 2023. doi: 10.48550/arXiv.2311.18154. arXiv: 2311.18154 [cs.R0].
- [111] I. Floris, J. M. Adam, P. A. Calderón, and S. Sales, “Fiber optic shape sensors: A comprehensive review,” *Optics and Lasers in Engineering*, vol. 139, 106508, 2021. doi: 10.1016/j.optlaseng.2020.106508.
- [112] K. C. Galloway, Y. Chen, E. Templeton, B. Rife, I. S. Godage, and E. J. Barth, “Fiber optic shape sensing for soft robotics,” *Soft Robotics*, vol. 6, no. 5, pp. 671–684, 2019. doi: 10.1089/soro.2018.0131.
- [113] N. Nazeer, “Fibre optic shape sensing and load monitoring of adaptive aerospace structures,” PhD thesis, Delft University of Technology, The Netherlands, 2022. doi: 10.4233/uuid:f77b0e21-338d-4edf-81d6-8a2bd400d71a.
- [114] *Principles of photogrammetry and remote sensing*, 2025. <https://uou.ac.in/sites/default/files/slm/GEOG-504.pdf> (visited on 15/5/2025).
- [115] Itseez, *Opencv: Open source computer vision library*, version 2.4, 2014. <https://opencv.org/> (visited on 20/5/2025).
- [116] Kamera Express, *GoPro HERO9 Black*, 2025. <https://www.kamera-express.nl/gopro-hero-9-black> (visited on 29/9/2025).
- [117] GoPro, Inc., *Gopro hero9 black user manual*, 2020. https://gopro.com/content/dam/help/hero9-black/manuals/HERO9Black_UM_ENG_REVA.pdf (visited on 1/12/2025).
- [118] OpenCV, *Camera calibration and 3d reconstruction*, 2020. https://docs.opencv.org/3.4/d4/d94/tutorial_camera_calibration.html (visited on 14/5/2025).
- [119] OpenCV, *Camera calibration and 3d reconstruction module (calib3d)*, 2023. https://docs.opencv.org/4.x/d9/d0c/group__calib3d.html (visited on 18/5/2025).
- [120] OpenCV, *Camera calibration using opencv (python)*, 2023. https://docs.opencv.org/4.x/dc/dbb/tutorial_py_calibration.html (visited on 20/5/2025).
- [121] T. Dang, C. Hoffmann, and C. Stiller, “Self-calibration for active automotive stereo vision,” in *2006 IEEE Intelligent Vehicles Symposium*, 2006, pp. 364–369. doi: 10.1109/IVS.2006.1689655.

- [122] D. Gallup, J.-M. Frahm, P. Mordohai, and M. Pollefeys, “Variable baseline/resolution stereo,” in *2008 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008, pp. 1–8. doi: 10.1109/CVPR.2008.4587671.
- [123] Roboception GmbH, *Rc-visard documentation: Stereo matching*, 2023. https://doc.rc-visard.com/v1.3/en/stereo_matching.html (visited on 14/4/2025).
- [124] H. Hirschmuller and S. Gehrig, “Stereo matching in the presence of sub-pixel calibration errors,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 437–444. doi: 10.1109/CVPR.2009.5206493.
- [125] T. Ewbank, “Efficient and precise stereoscopic vision for humanoid robots,” MSc thesis, Université de Liège, Belgium, 2017. <https://matheo.uliege.be/handle/2268.2/3144>.
- [126] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2003, ISBN: 9780511811685. doi: 10.1017/CBO9780511811685.
- [127] OpenCV, *Changing the contrast and brightness of an image*, 2025. https://docs.opencv.org/4.x/d3/dc1/tutorial_basic_linear_transform.html (visited on 3/9/2025).
- [128] OpenCV, *Contours: Getting started*, 2024. https://docs.opencv.org/3.4/d4/d73/tutorial_py_contours_begin.html (visited on 2/6/2025).
- [129] J. W. Kuhn, “The hungarian method for the assignment problem,” *Naval Research Logistics Quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955. doi: 10.1002/nav.3800020109.
- [130] C. Tomasi and T. Kanade, “Detection and tracking of point features,” School of Computer Science, Carnegie Mellon University, Tech. Rep. CMU-CS-91-132, 1991. <https://api.semanticscholar.org/CorpusID:238434334> (visited on 1/9/2025).
- [131] OpenCV, *Optical flow*, 2024. https://docs.opencv.org/4.x/d4/dee/tutorial_optical_flow.html (visited on 2/6/2025).
- [132] Pozyx Labs, *Pypozyx: Python library for the pozyx uwb system*, version 1.3.0, 2022. <https://github.com/pozyx-labs/pypozyx> (visited on 1/9/2023).
- [133] R. He, Z. Zhong, B. Ai, J. Ding, and K. Guan, “Analysis of the relation between fresnel zone and path loss exponent based on two-ray model,” *IEEE Antennas and Wireless Propagation Letters*, vol. 11, pp. 208–211, 2012. doi: 10.1109/LAWP.2012.2187270.
- [134] B. M. R. Franca, “A tool for aerodynamic analysis of flexible kites,” MSc thesis, Delft University of Technology, The Netherlands, 2014. <https://resolver.tudelft.nl/uuid:b972cc01-7a1e-449b-a4e2-503ee06adff6>.
- [135] Y. Gharbia, J. F. Derakhshandeh, M. M. Alam, and A. M. Amer, “Developments in wingtip vorticity mitigation techniques: A comprehensive review,” *Aerospace*, vol. 11, no. 1, 36, 2024. doi: 10.3390/aerospace11010036.
- [136] F. Gustafsson, “Determining the initial states in forward-backward filtering,” *IEEE Transactions on Signal Processing*, vol. 44, no. 4, pp. 988–992, 1996. doi: 10.1109/78.492552.
- [137] Massachusetts Institute of Technology, *6.s078 fine-grained algorithms and complexity: Lecture 17 — algorithms for finding long paths (part 1)*, 2020. <https://people.csail.mit.edu/virgi/6.s078/lecture17.pdf> (visited on 4/9/2025).
- [138] J. A. W. Poland, J. M. van Spronsen, M. Gaunaa, and R. Schmehl, “Wind tunnel load measurements of a leading-edge inflatable kite rigid scale model,” *Wind Energy Science Discussions*, 2025, [preprint]. <https://doi.org/10.5194/wes-2025-77>.
- [139] J. A. W. Poland, E. Fritz, and R. Schmehl, “Flow field analysis of a leading-edge inflatable kite rigid scale model using stereoscopic particle image velocimetry,” *Wind Energy Science Discussions*, 2025, [preprint]. <https://doi.org/10.5194/wes-2025-217>.
- [140] O. Cayon, “Fast aeroelastic model of a leading-edge inflatable kite,” MSc thesis, Delft University of Technology, The Netherlands, 2022. <https://resolver.tudelft.nl/uuid:aede2a25-4776-473a-8a75-fb6b17b1a690>.
- [141] O. Cayon, S. Watson, and R. Schmehl, “Kite as a sensor: Wind and state estimation in tethered flying systems,” *Wind Energy Science*, vol. 10, no. 10, pp. 2161–2188, 2025. doi: 10.5194/wes-10-2161-2025.
- [142] P. Roeleveld, “Fast finite element modelling of bridled leading-edge inflatable kites,” MSc thesis, Delft University of Technology, 2026.

- [143] W. Ciciora, J. Farmer, D. Large, and M. Adams, “Chapter 14 - linear microwave signal transportation,” in *Modern Cable Television Technology (Second Edition)*, W. Ciciora, J. Farmer, D. Large, and M. Adams, Eds., Morgan Kaufmann, 2004, pp. 593–616. doi: 10.1016/B978-155860828-3/50016-3.