


# Image processing for hydrofoil strut ventilation at moderate Froude number

A data driven approach to calmwater flow ventilation for vertical hydrofoil struts

AE5222: MSc. Thesis Aerospace Engineering  
Mees de Graaf

Delft University of Technology



# Image processing for hydrofoil strut ventilation at moderate Froude number

A data driven approach to calmwater flow  
ventilation for vertical hydrofoil struts

by

Mees de Graaf

Student Name	Student Number
Mees de Graaf	4867246

Supervisors: Dr. N.A.K. Doan, Dr. D. Fiscaletti, Dr. B. Font, Dr. G. Jacobi  
Substitute Supervisor: Dr. R.P. Dwight  
Course: MSc. Thesis Flight Performance  
Faculty: Faculty of Aerospace Engineering, Delft University of Technology  
Research group: Ship Hydromechanics, Department of Maritime Technology, Mechanical Engineering

# Preface

*At the closing stage of my time as a student in Delft, a chapter of my life closes with the completion of this thesis. Throughout my studies, I've had the incredible opportunity to discover my true passions and to meet many inspiring people along the way. My lifelong passion for sailing has not only shaped my personal life but has also been the driving force behind my academic journey. The moment I figured out that a sailboat is physically best described as an aircraft wing on its side, I knew exactly what I wanted to study. I have had no regrets since. This deep-seated interest culminated in the opportunity to research a subject that sits at the heart of high-performance sailing: The ventilation of surface-piercing hydrofoil struts.*

*My sincere gratitude goes to my diverse cohort of supervisors for their invaluable guidance. I am deeply grateful to Dr. Daniele Fiscaletti, Dr. Bernat Font, and Dr. Gunnar Jacobi from the Department of Ship Hydromechanics for sharing their expertise and helping me navigate the complexities of ventilated flows. I would also like to thank Dr. Anh Khoa Doan for his support as my initial supervisor, as well as Dr. Richard Dwight for stepping in as the substitute for Dr. Doan when he started a new chapter of his life working for University College London.*

*I would like to express a special thanks to the friendships that have carried me through this demanding study. Whether it is going to the faculty bar on a Friday afternoon, or a late-night study session at Echo, or an extraordinarily long lunch break. Those little moments are what made this journey something truly special. Lastly, I want to give a big thank you to my family for always supporting and encouraging me in my academic endeavors and in my sailing!*

Mees de Graaf  
Delft, June 2026

# Abstract

This research explores whether a quasi-automated framework for analysis and classification of flow features from visual data of a ventilated hydrofoil strut can be developed. Ventilated flows are flows in which air is entrained between the body of a normally wetted surface. They are particularly relevant for hydrofoiling boats, shallowly submerged ship propellers, superventilating torpedoes, and rudders on displacement and planing vessels. This research focuses on the natural ventilation of a hydrofoil strut at a yaw angle at a moderate Froude number. In such cases, air from the free surface flows towards the suction side of the hydrofoil.

A methodology was developed to extract physically meaningful features from high-speed imagery obtained during towing tank experiments. The framework combines unsupervised clustering techniques with image processing algorithms for geometric reconstruction.

It is demonstrated that the points on the hydrofoil free surface can be reconstructed algorithmically using an implementation of the Hough Gradient Method. These points have been successfully used to fit an asymmetric sigmoid function to the free surface contour. A modified implementation of the Hough Gradient Method, with adjustments to improve image contrast, enables reconstruction of the cavity closure angle. While unsupervised clustering methods show partial correspondence with known flow regimes, the results indicate that more advanced techniques are required for reliable, generalizable classification.

An additional contribution is the identification of an observable difference in free surface contour for a specific ventilation mechanism. Specifically, cases associated with a Laminar Separation Bubble exhibit a more forward free surface drawdown and a steeper depth gradient compared to other ventilation mechanisms. These findings demonstrate the potential of automated visual analysis not only for post-processing but also for advancing the physical understanding of hydrofoil ventilation.

# Contents

<b>Preface</b>	<b>i</b>
<b>Abstract</b>	<b>ii</b>
<b>Nomenclature</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Context and Motivation . . . . .	1
1.2 Background and Literature Review . . . . .	2
1.2.1 Flow Physics of Hydrofoil Ventilation . . . . .	2
1.2.2 Image Processing for Multiphase Flows . . . . .	8
1.2.3 Machine Learning in Flow Analysis . . . . .	11
1.3 Knowledge Gap and Research Objectives . . . . .	12
<b>2 Methodology</b>	<b>14</b>
2.1 Scope of the Dataset . . . . .	14
2.2 Practical Considerations and Data Pre-processing . . . . .	16
2.3 Flow Regime Classification . . . . .	18
2.4 Image Segmentation: Feature Extraction and Tracking . . . . .	24
<b>3 Results</b>	<b>38</b>
3.1 Flow Regime Classifier . . . . .	38
3.2 Free Surface Tracking . . . . .	44
3.3 Cavity Closure Angle . . . . .	46
<b>4 Conclusion</b>	<b>48</b>
4.1 Discussion . . . . .	51
4.1.1 Recommendations for Further Research . . . . .	52
<b>References</b>	<b>54</b>

# List of Figures

1.1	Explanation of terminology related to hydrofoil configurations, images from [25, 11] . . .	2
1.2	Ventilated cavities on a 2D hydrofoil section. (a) FW, (b) BV, (c) PV, (d) FV. Figure reproduced from [19] . . . . .	3
1.3	Illustrations of flow regimes on a 3D hydrofoil strut. Figures reproduced from [19] . . . .	4
1.4	Re-entrant jet on a simulated ventilated cavity at $F_{nh} = 1.75$ , $\alpha = 10^\circ$ . Figure reproduced from [18] . . . . .	5
1.5	Chord-based Froude number vs Angle of Attack map with inception boundaries for a submerged aspect ratio of 1. FW = Fully Wetted, FV = Fully ventilated, TVV = Tip Vortex Ventilation, VC = Vaporous Cavitation, VCV = Vaporous Cavity induced Ventilation. Figure reproduced from [22] . . . . .	6
1.6	Depth-based Froude number vs Angle of Attack map with inception events and annotations. Figure reproduced from [17] . . . . .	7
1.7	Illustration of ventilation inception mechanisms on a hydrofoil strut. Figure reproduced from [40] . . . . .	8
1.8	Underwater image of Partially Cavitating flow and a 1 second (500 frames) time average of pixel variance with isolines of pixel variance. Images reproduced from [17] . . . . .	9
1.9	Hough gradient method applied to horizon tracking from onboard footage in Olympic sailing. Figure reproduced from [32] . . . . .	11
2.1	Geometries of experimental test setup . . . . .	14
2.2	Camera views of experimental test setup . . . . .	15
2.3	Dataframe structure of reduced dataset in HDF5 format . . . . .	17
2.4	Example frames from the full run NACA-0010-34, $AR_h = 1.0$ , variable AoA, $F_{nh} = 2.5$ , Underwater straight on view, 100 Hz framerate . . . . .	18
2.5	Diagram showing interaction between different data sources to yield the reduced dataset	18
2.6	Examples of a frame from run 2098 at T - 0.02 s ( $F_n = 0.75$ , AR = 1.5, variable AoA), indicating that the hydrofoil tip is outside of the field of view. . . . .	19
2.7	Example frames for different flow regimes . . . . .	21
2.8	Clustering based on flattened image arrays with KMeans algorithm K = 4 . . . . .	22
2.9	Clustering based on histogram of the image with KMeans algorithm K = 2 . . . . .	23
2.10	Constant geometrical features of the NACA hydrofoil for AR = 1.0, image from run 2004 (AoA = $0^\circ$ , $F_n = 1.5$ , constant AoA) . . . . .	25
2.11	Step-by-step algorithm for free surface detection . . . . .	25
2.12	Optimal threshold from Otsu's method for background removal in FW images (total 2741 frames) . . . . .	27
2.13	Examples of a frame from run 2080 ( $F_n = 2.5$ , AR = 1.0, variable AoA) frames and masks with different background removal thresholds . . . . .	27
2.14	Example of canny edge detection applied to frame nr 2082 (T - 0.5 s) of run 2044 ( $F_n = 1.5$ , AR = 1.0, variable AoA). . . . .	29
2.15	Example of Probabilistic Hough Transform applied to frame nr 2082 (T - 0.5 s) of run 2044 ( $F_n = 1.5$ , AR = 1.0, variable AoA). . . . .	30
2.16	Example of combined and filtered houghlines from frame nr. 2082 (T - 0.5 s) of run 2044	33
2.17	Example of curve fitting of the detected points to frame nr 2082 (T - 0.5 s) of run 2044. Curve parameters found: $A = 1.124$ , $k = -6.170$ , $x_0 = 0.595$ , $C = 0.024$ . RMSE = 0.017	34
2.18	Step-by-step algorithm for cavity closure angle detection . . . . .	34
2.19	Example contrast maximization run 2044 ( $F_n = 1.5$ , AR = 1.0, variable AoA) at T + 0.4 s (frame nr 2172) . . . . .	35
2.20	Example of Probabilistic Hough Transform applied to frame 2172 of run nr 2044 . . . .	35
2.21	Example of bandpass filter [ $5^\circ$ , $45^\circ$ ] applied to frame 2172 of run nr 2044 . . . . .	36

---

2.22	Detected cavity closure angle of frame 2172 from run 2044. $\Phi = 18.47^\circ$ . . . . .	37
3.1	Cluster centers from flattened image clustering $k = 4$ . . . . .	39
3.2	Silhouette score vs relative margin for the $k = 4$ flattened image model, scores obtained from the samples in the validation set . . . . .	40
3.3	Histogram of the inertia contribution for the $k = 4$ flattened image model, obtained from the samples in the validation set . . . . .	40
3.4	Variation of number of clusters $k$ and the resulting Inertia . . . . .	41
3.5	Cluster centers from KMeans histogram clustering $k = 2$ . . . . .	42
3.6	Silhouette score vs relative margin for the $k = 2$ histogram model, scores obtained from the samples in the validation set . . . . .	42
3.7	Histogram of the inertia contribution for the $k = 2$ histogram model, obtained from the samples in the validation set . . . . .	43
3.8	Confusion Matrix obtained from the validation data using the $k = 2$ histogram clustering. . . . .	43
3.9	Histogram of the error and RMSE of the free surface curve fitting . . . . .	44
3.10	Summary of the success percentage of the Free Surface Reconstruction. . . . .	45
3.11	Per Ventilation mechanism mean asymmetric sigmoid curves for every frame before ventilation . . . . .	45
3.12	Per Ventilation mechanism mean of the parameters in the asymmetric sigmoid curve for every frame before ventilation . . . . .	46
3.13	Froude nr vs Cavity closure angle detected . . . . .	47
4.1	Cluster centers from KMeans histogram clustering $k = 2$ . . . . .	48
4.2	Confusion Matrix obtained from the validation data using the $k = 2$ histogram clustering. . . . .	49
4.3	Per Ventilation mechanism mean asymmetric sigmoid curves for every frame before ventilation . . . . .	50
4.4	Per Ventilation mechanism mean of the parameters in the asymmetric sigmoid curve for every frame before ventilation . . . . .	50
4.5	Froude nr vs Cavity closure angle detected . . . . .	51

# List of Tables

2.1	summary table of independent variables in the dataset . . . . .	16
2.2	Number of frames per flow regime in the pruned dataset . . . . .	20
2.3	Parameter values for Probabilistic Hough Transform used in Free surface tracking . . . . .	30
2.4	Parameter values for Probabilistic Hough Transform used in cavity closure angle . . . . .	36

# Nomenclature

## Abbreviations

Abbreviation	Definition
BV	Base Ventilation
CAD	Computer Aided Design
CFD	Computational Fluid Dynamics
CNN	Convolutional Neural Network
FV	Fully Ventilated flow regime
FW	Fully Wetted flow regime
PV	Partial Ventilation
PIV	Partical Image Velocimetry
SIFT	Scale Invariant Feature Transform

## Symbols

Symbol	Definition	Unit
$AR_h$	Submerged aspect ratio	[-]
$c$	Hydrofoil chord	[m]
$C_L$	3D Lift Coefficient	[-]
$C_D$	3D Drag Coefficient	[-]
$C_{M_{0.5c}}$	3D Moment Coefficient around half chord	[-]
$F_{nh}$	Depth based Froude number	[-]
$g$	Gravitational acceleration	[m/s <sup>2</sup> ]
$h$	Submergence depth	[m]
$p_c$	Pressure inside cavity	[Pa]
$p_\infty$	Freestream static pressure	[Pa]
$Re$	Reynolds number	[-]
$U_\infty$	Freestream velocity	[m/s]
$We$	Weber number	[-]
$z$	Submergence depth with respect to the undisturbed free surface	[-]
$\alpha$	Angle of Attack	[°]
$\mu_\infty$	Freestream viscosity	[Pas]
$\rho$	Density	[kg/m <sup>3</sup> ]
$\sigma$	Surface tension	[N/m]
$\sigma_c$	Cavitation Number with respect to atmospheric pressure	[-]
$\sigma_v$	Cavitation Number with respect to vapour pressure of water	[-]
$\Phi$	Cavity Closure angle	[°]

# Introduction

## 1.1. Context and Motivation

Over the past 15 years, hydrofoils have become a key technology in high-performance sailing, revolutionizing the top speeds achieved by modern racing yachts. The pioneering concepts from the foiling moth class and the last five editions of the America's Cup have now been widely adopted. Even round-the-world ocean racers, such as the Ultim class trimarans and the IMOCA 60 monohulls, are now equipped with hydrofoils [6, 33, 10].

However, the implementation of hydrofoils has not been without its challenges. Complex multiphase flow phenomena such as cavitation and ventilation limit the maximum hydrodynamic loading that the hydrofoils can produce, imposing limits on both speed and maneuverability [8, 40].

This thesis research focuses on flow ventilation, the entrainment of air under the normal free surface of an air-to-water interface. Ventilation is closely related to its more common and more widely researched cousin, cavitation. Cavitation is the entrainment of any gaseous medium (typically water vapour) in a normally liquid flow. Ventilation can lead to the formation of a ventilated cavity on the suction side of a hydrofoil. The ventilated cavity can span only a fraction of the chord length or extend beyond the trailing edge of the hydrofoil. It can be static in time or oscillate, and the air can be entrained under the free surface through various pathways [18, 16].

Ventilation is especially relevant for hydrofoil struts, surface-piercing hydrofoils, and rudders. Ventilation on any of these lifting surfaces can lead to two catastrophic outcomes: a loss of lift and a loss of lateral control, both of which can lead to splashdowns, nosedives, and capsizes of foiling sailboats. During these crashes, the crew and structure of the vessel are exposed to very high dynamic loads with the risk of structural failure and injury as a consequence [12].

Beyond its relevance for the sailing of hydrofoiling boats, the study of ventilation provides insight into the flow phenomenon from an isolated laboratory perspective. The findings from such research may expand our understanding of the design and optimization of shallowly submerged ship propellers, superventilating torpedoes, and rudders on displacement and planing vessels [40].

This thesis builds upon an experimental campaign by Manuel Ferreira and Carlos Rodrigues performed at the No.1 Towing Tank Facility of Delft University of Technology at the department of Maritime Engineering [4, 30]. They performed towing tank tests on various hydrofoil struts in a range of conditions.

From these experiments, high-resolution, high-frame-rate above-water and underwater videos are available. The motivation for making these videos was the lack of high-resolution videography in prior research on ventilation. In addition, these videos were made as a stepping stone to, at a later stage, make a tomographic Particle Image Velocimetry (PIV) reconstruction of the flow. The dataset is published on the 4TU research data platform [3].

Thus far, authors have identified flow regimes and flow features, such as the cavity closure line, free surface drawdown due to wavemaking, and initial air pockets at ventilation inception, by examining the flow. This approach has been sufficient because the experimental videography has been limited

in numbers. Of course, manually sifting through every frame of a towing tank campaign becomes a big hassle as the number of tests increases. In addition, the manual identification of flow regimes introduces the author's bias into the research. The need for an automated image processing approach is therefore evident. The central research question in this report will therefore be:

*How can a framework be developed for the quasi-automated analysis and classification of flow features from visual data of a ventilated hydrofoil strut?*

## 1.2. Background and Literature Review

This research combines the fields of multiphase flow physics, image processing, and machine learning in an attempt to answer the research question. Each field has its designated literature review section. Section 1.2.1 focuses on the flow physics of hydrofoil ventilation. Section 1.2.2 describes the current body of research related to image processing in experimental fluid dynamics. Finally, Section 1.2.3 describes relevant machine learning techniques for the problem.


### 1.2.1. Flow Physics of Hydrofoil Ventilation

Ventilation occurs when air is entrained into the low-pressure region of a hydrofoil strut or lifting surface, forming a cavity that alters the hydrodynamic loading. There are three conditions that need to be met before a flow might ventilate. Namely, a lower than atmospheric pressure somewhere on the surface of the hydrofoil, flow separation, and an air ingress path from the free surface to the surface of the foil [38, 35].

A hydrofoil strut serves as the structural connection between the hydrofoil wing and the hull, transferring the lift generated by the submerged foil to the hull. Its secondary function is to create a side force for lateral control and boat stability. A hydrofoil strut takes the shape of a slender appendage used to hold a hydrofoil wing underwater, creating a 'T-Foil'. The alternative geometry to a T-Foil would be a surface-piercing hydrofoil like a J-foil. The primary difference between a T-Foil and a J-Foil is the fact that the vessel's ride height is passively controlled with a J-foil, whereas a T-Foil requires a ride height control system [15]. Figure 1.1 illustrates the terminology used to describe different hydrofoil configurations. A hydrofoil strut interacts with both the free surface and the surrounding flow field, making it especially prone to ventilation [2].



(a) Bieker Moth sailing in foiling mode at upwind point of sail




(b) T-Foil configuration on a Bieker Moth CAD model



(c) J-Foil main foil with T-Foil rudder on a 69F sailboat

**Figure 1.1:** Explanation of terminology related to hydrofoil configurations, images from [25, 11]





Ventilation on a 2D hydrofoil section is classified into 3 distinct flow regimes. Namely, Fully Wetted (FW) flow, Partially Ventilated (PV) flow, and Fully Ventilated (FV) flow. Furthermore, on blunt bodies or hydrofoils with a blunt trailing edge, an additional mode of base ventilation (BV) exists. Base ventilation occurs in the wake of the blunt body. The streamlines for these different flow regimes are illustrated in Figure 1.2.



**Figure 1.2:** Ventilated cavities on a 2D hydrofoil section. (a) FW, (b) BV, (c) PV, (d) FV. Figure reproduced from [19]

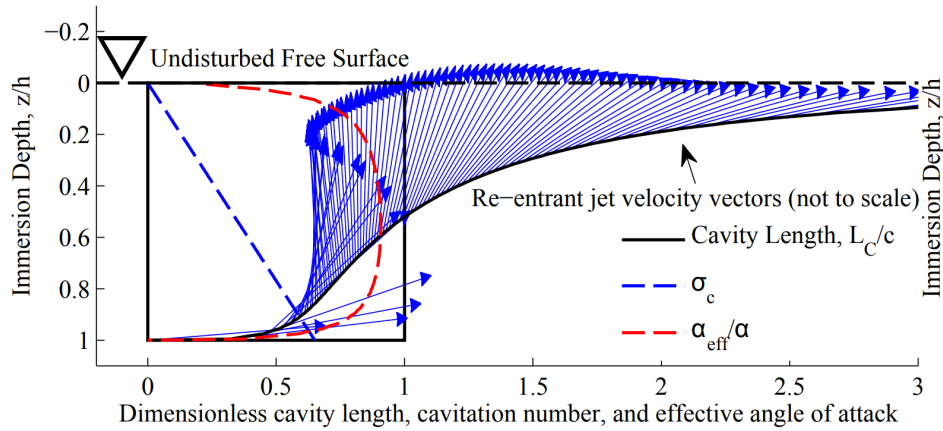
When changing from a 2D hydrofoil section to a 3D hydrofoil strut, three important geometric features of the flow determine the cavity topology. They are in no particular order: The cavity closure line, the free surface deformation, and the spray sheet. Harwood, Young, and Ceccio noted the importance of a clearer definition of ventilation regimes on a 3D hydrofoil. The features are illustrated on top of towing tank images by Harwood et al. in Figure 1.3 [19].

On a 3D hydrofoil strut, there exists a depth-wise variation in the topology of the ventilated cavity. There are again three flow regimes denoted: Fully Wetted (FW), Partially Ventilated (PV), and Fully Ventilated (FV). A fully wetted hydrofoil strut causes only wavemaking displacement of the free surface and a spray sheet, but not a clearly definable cavity (Figure 1.3a). A partially ventilated hydrofoil strut causes a ventilated cavity on its suction side that extends to a depth smaller than its total submergence depth (mode 1 Figure 1.3c). Alternatively, if the mean angle of the cavity closure line is larger than  $45^\circ$ , this is also considered partial ventilation (mode 2 Figure 1.3d). On a partially ventilated hydrofoil strut, the majority of the cavity extends beyond the chord length such that a cross-section at a shallow depth would be classified as 2D full ventilation. Finally, a fully ventilated hydrofoil strut leaves a cavity that extends to the full submergence depth and has a mean cavity closure angle smaller than  $45^\circ$  with respect to the horizon (Figure 1.3b).

(a) FW  $F_{nh} = 2, \alpha = 10^\circ$ (b) FV  $F_{nh} = 3, \alpha = 15^\circ$ (c) PV mode 1  $F_{nh} = 1, \alpha = 20^\circ$ (d) PV mode 2  $F_{nh} = 1.54, \alpha = 7.5^\circ$ **Figure 1.3:** Illustrations of flow regimes on a 3D hydrofoil strut. Figures reproduced from [19]

Young and Harwood describe an additional topological feature of ventilated flows that is particularly relevant for ventilation on 3D hydrofoil struts. Namely, a Re-entrant jet. The re-entrant jet originates from conservation of momentum around the cavity closure line. The incoming flow is redirected at an oblique angle due to spanwise flow effects. The result is an observable spray that extends upwards at an angle to the horizon, double that of the cavity closure line. The re-entrant jet is a key feature in the elimination and washout of ventilated cavities, as it provides a path for water to re-flood the

cavity. Young and Harwood both note that the angle of the cavity closure line, and with it the angle of the re-entrant jet, is a key criterion for determining the stability of a ventilated cavity [18, 40]. The re-entrant jet is illustrated in Figure 1.4.



**Figure 1.4:** Re-entrant jet on a simulated ventilated cavity at  $F_{n_h} = 1.75$ ,  $\alpha = 10^\circ$ . Figure reproduced from [18]

In an attempt to reduce the scale of the ventilated flow problem, numerous authors have performed dimensional analyses [38, 16, 40]. The following non-dimensional numbers are of primary importance;

- Depth-based Froude Number:

$$F_{n_h} = \frac{U_\infty}{\sqrt{gh}} \quad (1.1)$$

- Submerged aspect ratio:

$$AR_h = \frac{h}{c} \quad (1.2)$$

- Cavitation number:

$$\sigma_c = \frac{p_\infty - p_c}{\frac{1}{2}\rho U_\infty^2} \quad (1.3)$$

In addition to (1.1), (1.2), and (1.3), the angle of attack  $\alpha$  is also crucial in the physics of ventilated flows. Note that the angle of attack is sometimes referred to as a yaw angle, given the typical orientation of a hydrofoil strut on a vessel.

In addition, note that for a naturally ventilated cavity, the pressure inside the cavity is equal to atmospheric pressure, and the freestream static pressure is the sum of atmospheric pressure and hydrostatic pressure ( $p_c = p_{atm}$ , and  $p_\infty = p_{atm} + \rho gz$ ). Substituting into (1.3) and recognizing the inverse square of the Froude number results in (1.4). The cavitation number is therefore linearly increasing with submergence depth along the hydrofoil strut [39].

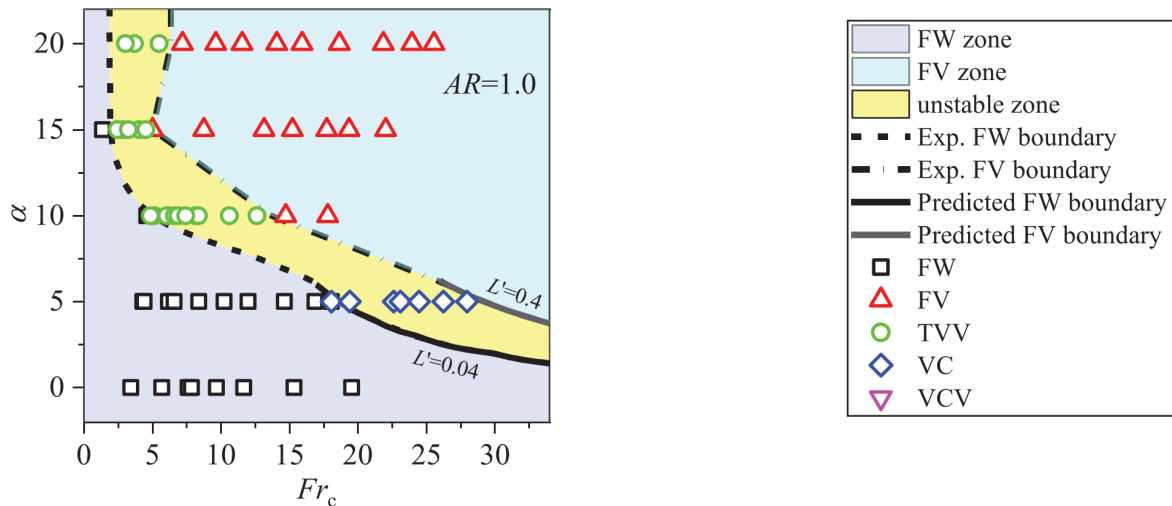
$$\sigma_c = \frac{\rho gz}{\frac{1}{2}\rho U_\infty^2} = \frac{1}{2F_{n_h}^2} \frac{z}{h} \quad (1.4)$$

The change in non-dimensional force coefficients  $C_L$ ,  $C_D$  and  $C_{M_{0.5c}}$  due to a ventilated cavity has been studied by numerous authors [1, 9, 39]. The general conclusion is that ventilation causes a loss of lift, a reduction in friction drag that is compensated by the addition of spray drag and form drag, and a decrease in moment coefficient around mid-chord.

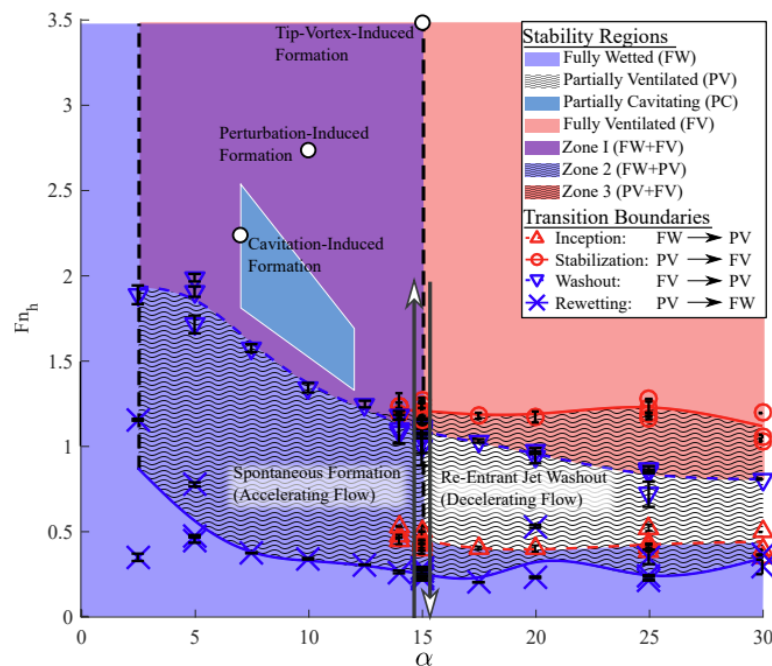
In addition to the aforementioned numbers, the Reynolds number and Weber Number, which describe the relative importance of viscosity and surface tension with respect to inertia in the flow, respectively, also play a secondary role. The Reynolds number is defined as  $Re = \frac{\rho U_\infty c}{\mu_\infty}$ , and the Weber number as  $We = \frac{\rho U_\infty^2 c}{\sigma}$ . The Reynolds number is critical in flow separation of a foil section. Wetzel found

surface tension unimportant for Weber numbers  $We \geq 250$  [38]. Weber numbers in the order  $O(10^4)$  are commonly achieved in towing tank research, and at vessel scale this will increase even further to  $O(10^5)$ .

Understanding when ventilation occurs is currently encapsulated in a map of inception conditions in Froude Number vs angle-of-attack space. These maps have been obtained from observations made during towing tank tests. In addition, there exist theoretical and semi-empirical boundaries for different types of ventilation inception. In this context, 'inception' refers to the transition from fully wetted to ventilated flow, whether PV or FV. Huang et al. focused their research on high Froude numbers, where the presence of vapour cavities is more likely Figure 1.5. Harwood drew a map at moderate Froude numbers based on his towing tank research; this can be seen in Figure 1.6.



**Figure 1.5:** Chord-based Froude number vs Angle of Attack map with inception boundaries for a submerged aspect ratio of 1. FW = Fully Wetted, FV = Fully ventilated, TVV = Tip Vortex Ventilation, VC = Vaporous Cavitation, VCV = Vaporous Cavity induced Ventilation. Figure reproduced from [22]

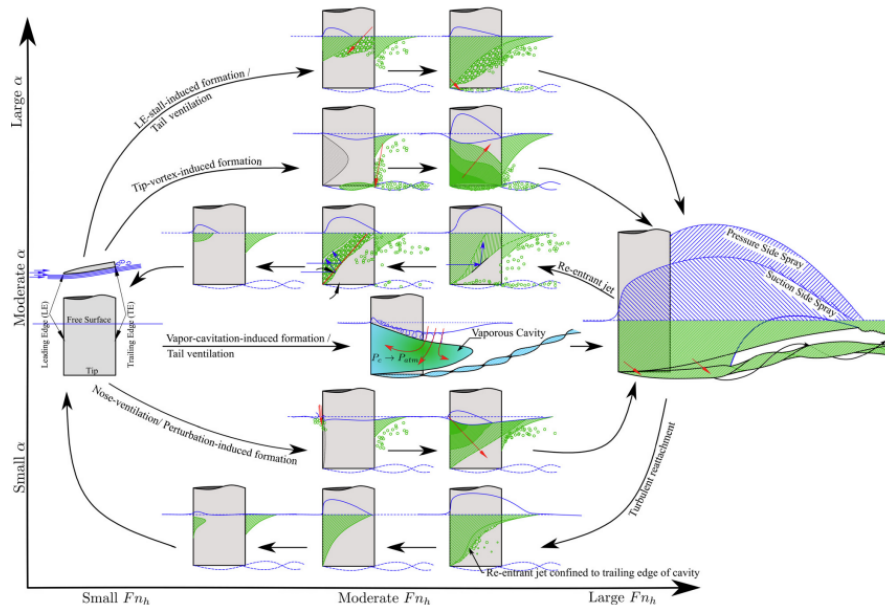


**Figure 1.6:** Depth-based Froude number vs Angle of Attack map with inception events and annotations. Figure reproduced from [17]

Both maps show a region where both FW, PV, and FV are possible. This region is referred to as 'bistable'. On the lift polar, this region causes an hysteresis loop [19]. Hysteresis refers to the fact that recovery of the ventilated cavity from PV or FV to FW occurs at a lower angle of attack or Froude number than at inception. Note that both Harwood and Huang et al. conducted their towing tank experiments at a constant Angle of Attack by accelerating to a Froude number. Ferreira and Rodriguez investigated the effect of varying the Angle of Attack at constant Froude number and found that the inception boundaries differ from Harwood [4].

According to Harwood, there are five distinct mechanisms by which a ventilated cavity can form on a hydrofoil strut. Each of them is associated with a different path of air ingress. They are illustrated in Figure 1.7. Note that these ventilation mechanisms refer to a cavity on the low-pressure side of a body and therefore exclude base ventilation.

1. Free surface perturbation
2. Laminar Separation Bubble
3. Rayleigh Taylor Instability
4. Tip vortex ventilation
5. Cavitation induced ventilation



**Figure 1.7:** Illustration of ventilation inception mechanisms on a hydrofoil strut. Figure reproduced from [40]

Carefully note that Ferreira and Rodrigues use a different nomenclature to refer to similar phenomena. They distinguish three distinct trigger mechanisms, namely Nose ventilation, Tail ventilation, and Base ventilation. They also observe that multiple trigger mechanisms can occur simultaneously during the inception of ventilation. Finally, they elaborate that for Nose ventilation, the existence of a Laminar Separation Bubble (LSB) is relevant. An LSB is a leading edge flow separation followed by turbulent reattachment. At varying angles of attack, the pressure distribution over the hydrofoil changes, and thus the LSB changes position and chordwise length. Ferreira hypothesizes that the absence of an LSB on thick hydrofoil struts and very sharp nosed hydrofoil struts (those exhibiting thin airfoil stall at the leading edge) results in a lower likelihood of Nose Ventilation occurring on these struts.

Regarding the inception of ventilation, there is another notable classification that should be noted. Wetzel identified two time scales at which ventilation might develop. A rapid, sudden onset vs a gradual increase of the ventilated cavity. Wetzel referred to these, respectively, as 'flash' and 'creeping' inception [38].

Thus far, the classification of ventilated flows has been a manual and qualitative exercise. The aforementioned studies distinguish between flow regimes based on observations made during the towing tank research. An automated way to classify these types of flows would greatly benefit further experimental research on hydrofoil ventilation. An automated classifier should, in principle, encapsulate the physical understanding of flow ventilation. More on the classification of flows in Section 1.2.2.


The inception and development of ventilation are further complicated by unsteady effects, nonuniform pressure distributions from the vessel hull or other appendages, and interactions with surface waves [20, 21, 41]. Understanding these mechanisms is deemed critical for a complete understanding hydrofoils but is considered beyond the scope of this research.

### 1.2.2. Image Processing for Multiphase Flows

Accurate analysis of ventilated flows relies heavily on the reliable identification of air, water, and solid surfaces. In previous studies on ventilation, flow features and phase boundaries have mostly been manually identified and drawn by the author. Image processing can play a key role in identifying the phase boundaries and flow features automatically. This section reviews image processing approaches used in multiphase flow studies and discusses their potential applications for flow ventilation on

hydrofoil struts.

In a subsection of his PhD thesis, Harwood described a technique for identifying unsteady regions in a partially cavitating flow from video footage. The technique is illustrated in Figure 1.8. The technique is simple but elegant. Unsteady regions of the flow are identified by averaging the variance in pixel intensity over multiple frames and then constructing contours on which this quantity is constant [17].



**Figure 1.8:** Underwater image of Partially Cavitating flow and a 1 second (500 frames) time average of pixel variance with isolines of pixel variance. Images reproduced from [17]

In 2023, Jacobi used the variance in pixel intensity for the automatic masking of PIV images [23]. In general, in PIV it is important to mask shadows, foil surfaces, and other distortions.

In 2025, Stichter et al. applied image processing techniques to Interferometric Particle Imaging (IPI) frames. The objective of the study was to measure the concentration and size of microbubbles in water. The IPI images were first blurred using a Gaussian blur, then a thresholding operation was applied, followed by the Canny edge detection algorithm and Hough transform. In addition, the 2D Fast Fourier Transform (FFT) was used to detect the fringe pattern within the microbubbles.


Franzosi et al. in 2023 performed a study on root cavitation erosion. Typically, the analysis of root cavitation is performed qualitatively. Franzosi et al. developed a data processing pipeline to automatically track cavitation bubbles from a high-speed video of a propeller. The method consisted of a masking step, in which the region of interest, namely the propeller root, was separated from other regions subject to cavitation, such as the tip and the suction side of the blade. Afterwards, the Scale Invariant Feature Transform (SIFT) algorithm was applied to extract anything interesting from the unmasked area. Three sequential filtering steps for outlier detection were used afterwards.

Finally, three techniques were used to estimate the size of the cavitated regions for each frame of the video. The three techniques were: Convex Polygon Hull, Feature Expansion, and Grayscale Pixel thresholding. The Convex Polygon hull was deemed insufficiently accurate, as it would highlight too large regions based on remaining outliers. Feature Expansion and Thresholding were both successful at highlighting cavitation regions after parameter tuning was applied. Franzosi et al. stressed the importance of lighting conditions in order to accurately depict the smallest bubbles in the video [14].

When using image processing to analyze hydrofoil ventilation, it is crucial to recognize that camera footage provides a restricted and potentially misleading representation of reality. A camera inherently introduces radial and tangential optical distortions. Depending on the viewing angle, lens characteristics, and location in the field of view, these distortions can be negligible or very significant. Refraction at the air-water interface is especially relevant for multiphase flows. Points on the hydrofoil surface may appear warped if a ventilated cavity is between the point of interest and the camera. Lighting conditions and lens glare can also alter the perceived appearance of multiphase flow features, making bubbles, cavities, or ventilation structures appear different from their true form. These factors highlight the need for advanced image pre-processing to ensure that image-based analyses reflect the underlying physical phenomena as accurately as possible.

In 2025, Franzosi et al. performed a study on bubble cavitation on a NACA 0015 hydrofoil. In this study, a comparison was made between the computer vision techniques active triangulation and shape-by-silhouette. Both of these techniques use at least 2 different camera angles. The goal of the study was to create a 3D time-resolved reconstruction of cavitation bubbles. A pinhole camera model was used for the shape-by-silhouette technique. In addition, the importance of accounting for camera lens distortions and air-to-water and plexiglass-to-water interfaces for accurate 3D reconstruction was stressed. In addition, the importance of optical access for multiple camera views was emphasized [13].

The detection of edges is deemed a promising technique for identifying cavity closure lines, and potentially masking/subtracting the grid that is drawn on the hydrofoil strut surface. Schutt successfully used a combination of the Canny edge detection algorithm and the Hough transform to track the horizon and draft stripes (taped lines of contrasting colour used to visualize the curvature of a sail) from video images of Olympic-level sailing. His research objective was to understand the kinetic propulsion techniques (accelerating the boat through bodyweight movements) learned by sailors from a fluid dynamics point of view. Figure 1.9 shows the algorithm applied to the tracking of the horizon from the onboard bow camera, whereby Schutt is able to verify the boat's heel angle [32].



**Figure 1.9:** Hough gradient method applied to horizon tracking from onboard footage in Olympic sailing. Figure reproduced from [32]

In their 2025 study on measuring the size of microbubbles suspended in water, Stigter et al. aimed to improve the quantitative reliability of Interferometric Particle Imaging (IPI) for measuring microbubble size distributions and concentrations. The research objective was to develop an improved calibration methodology and to provide a rigorous uncertainty assessment for microbubble concentration measurements. Their results are validated against Long-range Microscope Shadowgraphy (LMS), demonstrating agreement within uncertainty bounds.

In the image analysis stage of the IPI images, the authors employ the Hough Gradient Method to detect circular microbubbles in the IPI images. Prior to applying the method, each image undergoes background subtraction and Gaussian smoothing with a  $5 \times 5$  px kernel to suppress noise and high-frequency components. A minimum intensity threshold is then applied to isolate candidate bubble regions, after which edge detection is performed. The Hough gradient method uses these detected edges, along with local gradient information, to determine the position and area of circular bubble images. An initial diameter estimate is derived from the detected area, assuming circular geometry. Because this estimate is sensitive to threshold selection and overlapping bubbles, the authors refine the diameter measurement by extracting a cross-section through the bubble image parallel to the fringe pattern [34].

### 1.2.3. Machine Learning in Flow Analysis

The complexity of the ventilation phenomenon, combined with limited experimental validation data, poses a significant challenge for numerical modeling approaches (CFD) [37, 27]. Prior research has shown the difficulty of accurately predicting the inception and evolution of ventilation [19]. Machine learning can offer a promising alternative to traditional physical modeling by learning patterns directly from data. With the high-resolution video dataset from the experimental campaign of Ferreira and Rodrigues [4], machine learning can automate the classification of flow regimes and potentially predict the inception mechanism of ventilation.

Machine learning is increasingly used in fluid dynamics to analyze complex, time-dependent flow phenomena. Unlike traditional numerical simulations, which solve governing equations directly, machine learning models identify patterns in data, offering a computationally efficient alternative for certain tasks. In fluid dynamics, machine learning can serve as a potential solution for two types of problems, namely, classification and prediction. The classification tasks involve categorizing flow states. This report would entail distinguishing ventilated and non-ventilated regimes. The prediction task aims to forecast future flow behavior based on past observations. While machine learning can significantly reduce computational costs with respect to numerical simulations, training models requires large datasets and careful consideration of model generalization. This section reviews studies that have applied machine learning techniques to multiphase flow and discusses their potential relevance to hydrofoil ventilation research.

In 2025, Liu et al. performed a study on sheet cavitation on a NACA 0015 and NACA 4412 hydrofoil. A U-Net Convolutional Neural Network was used to segment cross-sectional images into airfoil and cavitation cloud or sheet. After segmentation, the algorithm determined the cavitation length and chord length in pixels to arrive at the numerical end result of the cavitation ratio  $l/c$  [24].

On the topic of classification, Nji et al in 2022 tested three unsupervised clustering algorithms, DBScan, Hierarchical Agglomerative Clustering (HAC), and KMeans, on spatio-temporal weather data from the ERA5 reanalysis dataset. The objective was to figure out which clustering algorithm has the best performance on the dataset. In addition, the clustering was preceded by a Principal Component Analysis (PCA), in order to study the effect of feeding differing numbers of PCA's to the clustering algorithm. The study provides a detailed explanation of the performance metrics employed to evaluate the clustering algorithms, including RMSE, Silhouette Score, and both intra- and inter-cluster correlation. It was found that Agglomerative Clustering and KMeans perform comparable [28].

### 1.3. Knowledge Gap and Research Objectives

Having explored the current body of research on hydrofoil strut ventilation and image postprocessing in experimental fluid dynamics, the central research question can be broken down into several sub-questions that serve as a benchmark for the framework for the quasi-automated processing of towing tank images for ventilation on hydrofoil struts.

In addition, a breakdown of the task can be made into three distinct sub-tasks, namely: Image pre-processing, Image classification, and image segmentation. This task breakdown serves as a skeleton for the modular design of an image processing pipeline.

Image pre-processing is any operation on the arrays of pixels, or on the metadata, in order to prepare the data for the other steps of the pipeline. Pre-processing starts with practical considerations from such a large dataset: Data selection for a balanced dataset, considerations of memory capacity, and the option to downsample. The practical considerations, as well as the scope of the dataset, are further elaborated on in Section 2.2.

Image classification is the automated categorization of images into different subgroups. It is important to note the distinction in nomenclature between the Image classification task and the typical definition of classification vs clustering in machine learning. Clustering in machine learning is an operation performed on unlabeled data where similarities are sought after in order to determine which data points belong together, whereas classification is typically performed on labeled data where the algorithm 'learns' to predict the class of new images based on the prescribed labels and the training data. This distinction is further elaborated on in Chapter 2, but in general, a mix of clustering and classification is used in order to achieve the image processing framework. In the framework, the primary objective of classification is to bin images into 4 categories, namely FW, FV, PV, or Transitional. It is also interesting to explore if an algorithm can distinguish more categories without overfitting.

Finally, image segmentation is the task of extracting regions of interest from a set of images based

on the numerical data embedded in the pixels. This is the task for the algorithms explored in subsection 1.2.2. For the image processing framework, the objective is to extract the contour of the free surface pre-ventilation, and cavity closure angle post-ventilation.

The subtasks of the image processing framework are further broken down with the following research subquestions:

- To what extent can a KMeans clustering classify the towing tank videography into categories related to their flow state; FW, Transitional, PV or FV?
- How robust is a quasi-automated framework to different foil geometries, wave conditions, and other camera angles?
- How well does a quasi-automated framework capture the human interpretation of flow ventilation?
- Can spatio-temporal features of the flow be reliably tracked?
- What image pre-processing steps are required to reliably extract flow features instead of image artifacts?
- What initial advancements can be made towards an algorithm to predict the inception of ventilation?


# Methodology

## 2.1. Scope of the Dataset

Before delving into image pre-processing and analysis, it is important to establish a foundational understanding of the dataset from both a fluid dynamics and imaging perspectives.

The experimental campaign conducted by Ferreira & Rodriguez focused on the study of ventilation phenomena around hydrofoils. four main geometries were investigated: Two different profiles and two aspect ratios. The profiles are a semi-ogival hydrofoil and a modified NACA 0010 hydrofoil. Their geometries are outlined in Figure 2.1. The submerged aspect ratios as defined in (1.2) are respectively 1.0 and 1.5. The experiments were carried out across a range of depth-based Froude numbers from 0.5 to 2.5, with all runs conducted at a chord-based Reynolds number higher than 500,000. The flow regimes observed in the dataset include Fully Wetted flow (FW), Partial Ventilation (PV), Full Ventilation (FV), transition from FW to PV and FW to FV, as well as washout from FV to FW and PV to FW. The flow regimes considered in this dataset correspond to distinct physical states of the air–water interface around the hydrofoil. Fully Wetted (FW) flow denotes the condition in which the hydrofoil surface remains entirely submerged and in contact with water, without air entrainment. Partial Ventilation (PV) refers to a state in which an air cavity forms locally, originating from the free surface and attaching to part of the hydrofoil surface while other regions remain wetted. Full Ventilation (FV) describes the situation in which the cavity extends over the entire suction side of the hydrofoil, resulting in a continuous air layer separating the foil from the surrounding water. Transitions from FW to PV or FV, therefore, represent the onset of ventilation, characterized by the growth and stabilization of an air cavity along the hydrofoil surface. A more detailed discussion of the physical mechanisms underlying these regimes is provided in Section 1.2.1, and a representative frame of the aforementioned flow regimes is shown in Figure 2.7.

It was decided to reduce the scope of this research to only the NACA 0010-34 geometry in calmwater conditions, only the transition from FW to PV or FV, and only the underwater perpendicular view. Thereby omitting base ventilation, washout mechanisms, and clear imagery of the spray sheet.




**Figure 2.1:** Geometries of experimental test setup

Each experimental run is identified by a unique four-digit run number. The initial digit signifies the configuration:

- **1XXX**: Semi-ogival hydrofoil in calm water,
- **2XXX**: NACA 0010-34 hydrofoil in calm water,
- **3XXX**: NACA 0010-34 hydrofoil in wave conditions.

During the tests, 4 cameras were used to capture video of the flow: 2 above the water surface and 2 underwater. Each camera has a different viewing angle to the foil surface. One is oriented perpendicular to the flow direction and positioned so that the suction side of the hydrofoil is in the middle of the field of view. The other is looking at an oblique angle from behind towards the suction side of the hydrofoil. The field of view for each camera is illustrated in Figure 2.2.



**Figure 2.2:** Camera views of experimental test setup

There is an Excel sheet with metadata from the towing tank test campaign. The experimental test matrix was defined using five key variables, comprising two continuous parameters and three categorical parameters.

The continuous parameters were the submerged aspect ratio, tested at two discrete levels, and the Froude number, tested at nine discrete levels.

Three categorical variables were included in the test design. First, two test execution modes were considered: (i) constant angle of attack at varying Froude number, and (ii) varying angle of attack at constant Froude number. Second, two foil cross-section geometries were tested, the sections outlined in Figure 2.1. Finally, tests were conducted under two environmental conditions: calm water and wave conditions. The variables and their levels are summarized in Table 2.1.

**Table 2.1:** summary table of independent variables in the dataset

<i>Variable</i>	<i>Levels</i>
Submerged Aspect Ratio	1.0, 1.5
Froude Number	0.5, 0.75, ... , 2.5
Angle-of-Attack Test Mode	Constant $F_{nh}$ , Constant $\alpha$
Foil Cross-Section	NACA 0010-34, Semi-Ogive
Free Surface Condition	Calm Water, Waves

In addition to these independent variables chosen by Ferreira and Rodrigues for their test matrix, the metadata also includes the following dependent variables [4]:

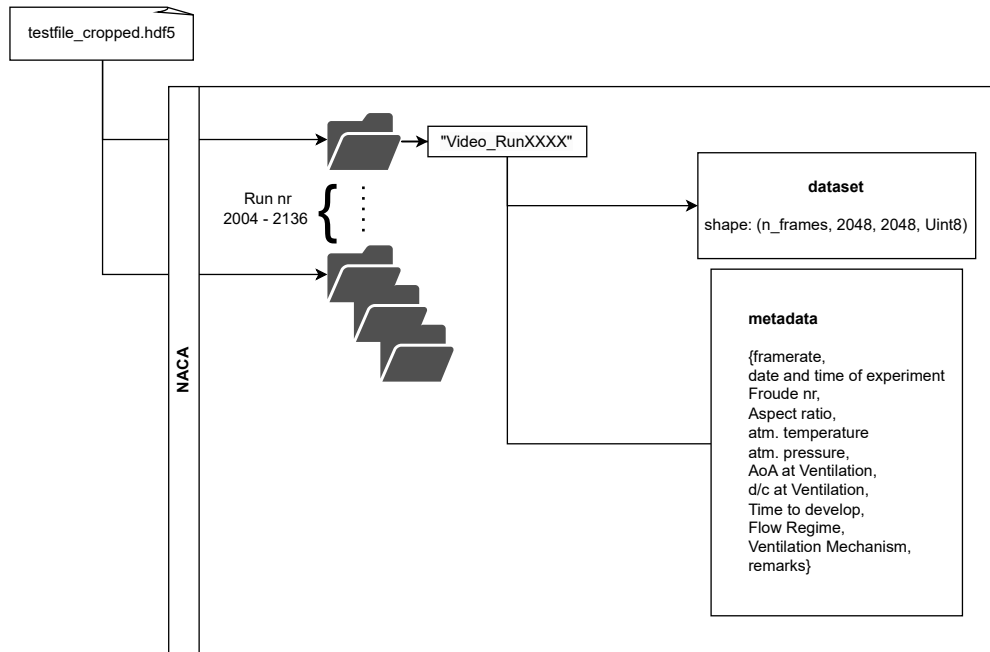
- Flow Regime {FW, PV, FV}
- Frame number of ventilation inception
- Time to develop a ventilated cavity
- Inception Mechanism {Nose Ventilation, Laminar Separation Bubble, Stall-Induced Ventilation, Tail Ventilation, Nose & Tail Ventilation}

Having established what is in the dataset allows for filtering and comparative analysis across flow conditions and geometrical configurations. Especially relevant later is the categorization into different ventilation inception mechanisms [4, 30]. Note that the inception mechanisms observed by Ferreira and Rodrigues in the dataset differ from the mechanisms observed and illustrated by Harwood and Young in Figure 1.7.

## 2.2. Practical Considerations and Data Pre-processing

The raw dataset comprises high-resolution video recordings stored in a mix of file formats: DaVis .ims sets from the underwater cameras, MP4 files from the above-water cameras, and Pickle files containing full image arrays from the MP4 files. The complete dataset spans approximately 10 TB of memory. Given the large volume of data, pre-processing steps are critical for improving the accessibility and usability of the vast dataset.


For usability and memory allocation purposes, it was decided to create a substitute database from the original in the HDF5 data format. As mentioned in Section 2.1, the substitute database only contains images from the underwater perpendicular view (Figure 2.2c) of the NACA 0010-34 geometry in calmwater conditions. HDF5 is a database format that supports partial I/O to bulk storage devices, as well as a convenient way to store metadata along with video images in the same database [36]. For extra fast I/O, every frame is stored uncompressed as a 2048 x 2048 array of 8-bit unsigned integers. The structure of the database is illustrated in Figure 2.3. The substitute database has a highly modular design with the capability for expansion with data from the other cameras and other test variables.



**Figure 2.3:** Dataframe structure of reduced dataset in HDF5 format

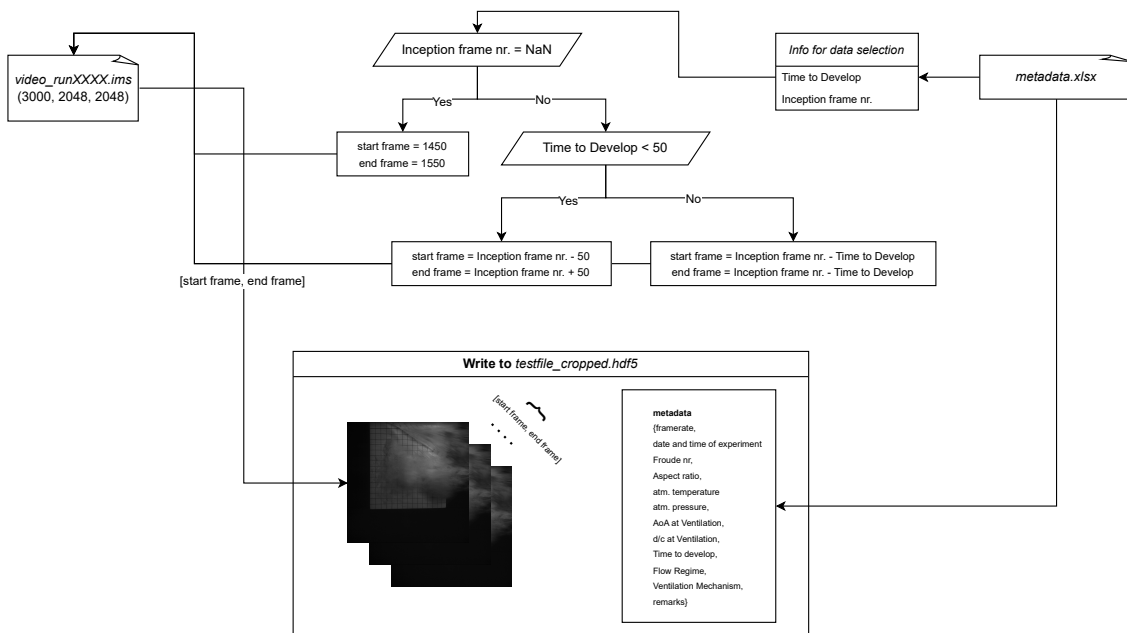
## Data Selection

A balanced dataset for model training is key to success, the current full dataset consists of relatively many images of a FW flow with a low drawdown ratio, superseded by a slow increase of the drawdown ratio. It can take, depending on the run number, as long as 2000 frames before ventilation inception. Methodical data selection can help minimize the number of redundant frames needed for processing and improve the training time of models. The frames around the inception of ventilation are the most interesting from an analysis point of view. Therefore, the runs in the reduced dataset will be pruned based on the inception frame and the time to develop a ventilated cavity. An example of the number of redundant frames is given in Figure 2.4. The information about the time to develop a ventilated cavity and the frame in which inception starts is available from the metadata Excel file of the original database. Figure 2.5 shows a visual diagram of the exact steps to select data for the reduced dataset.



**Figure 2.4:** Example frames from the full run NACA-0010-34,  $AR_h = 1.0$ , variable AoA,  $F_{n,h} = 2.5$ , Underwater straight on view, 100 Hz framerate

Figure 2.4b shows that from frame 0 to 1000, almost nothing changes. But it does take 1000 image arrays  $\approx 3.9$  Gb of memory. This was the primary motivation to reduce the dataset.



**Figure 2.5:** Diagram showing interaction between different data sources to yield the reduced dataset

## 2.3. Flow Regime Classification

Because the ventilation phenomena manifest themselves as distinct flow regimes, each characterized by specific geometric and optical features in the image data, an automatic method for assigning frames


to the appropriate processing pipeline is desirable. Given that the different physical features of a ventilated cavity occur on different parts of the hydrofoil and have very different geometries, it was deemed realistic to develop multiple feature extraction and tracking algorithms, each corresponding to a dominant feature for a given flow regime.

The objective of this section is to determine whether the image data naturally separates into clusters that correspond to physically meaningful flow regimes, enabling automated and consistent regime identification. The following sections describe the construction of a balanced dataset, the algorithm used for clustering, and the metrics used to evaluate the resulting clustering.

### Data Exploration and Pre-processing

To distinguish between different ventilation regimes (FW, transitional, PV, and FV), an unsupervised clustering approach is adopted. Although flow regime labels are available for a subset of the dataset (Ferreira and Rodrigues), transitional frames are not explicitly labeled. The clustering procedure is therefore used to determine whether natural groupings in the image data correspond to physically observable flow regimes, without explicitly training on regime labels.

As an initial step, the content of the pruned dataset was examined using the available regime annotations. The analysis was restricted to the NACA-0010-34 hydrofoil with aspect ratio  $AR = 1.0$ . For the  $AR = 1.5$  configuration, the hydrofoil tip lies partially outside the camera field of view, as illustrated in Figure 2.6. Since the tip region is very relevant for fully ventilated flows, these cases were excluded to ensure consistent spatial information across samples.



**Figure 2.6:** Examples of a frame from run 2098 at  $T = 0.02$  s ( $F_n = 0.75$ ,  $AR = 1.5$ , variable AoA), indicating that the hydrofoil tip is outside of the field of view.

An inspection of the labeled data revealed a significant class imbalance. The Fully Wetted (FW) regime is overrepresented compared to the ventilated regimes. The distribution of frames per regime is summarized in Table 2.2. In particular, FW contains approximately 14 times more samples than FV.

**Table 2.2:** Number of frames per flow regime in the pruned dataset

<i>Flow Regime</i>	<i>Number of Frames</i>
FW	2401
PV	203
FV	170
Transitional	1428

Such an imbalance would bias the clustering process, as centroid locations in KMeans are influenced by sample density. To mitigate this effect, a stratified sampling procedure was applied to construct a balanced subset.

The sampling strategy consisted of:

- Partitioning the dataset into containers based on the available regime labels (FW, PV, FV).
- Constructing a transitional container containing frames between the inception frame ( $n_{\text{frames}}/2$ ) and the full cavity frame.
- Drawing random samples within each container on a per-run basis to preserve temporal diversity.

The smallest regime (FV) contains 170 frames; therefore, the maximum fully balanced dataset consists of  $4 \times 170 = 680$  frames. To enable independent evaluation of clustering performance, part of the data was withheld for validation.

The validation set consists of two complete runs: one labeled FV and one labeled PV. Together, these runs span the full progression of regimes (FW  $\rightarrow$  transitional  $\rightarrow$  PV and FW  $\rightarrow$  transitional  $\rightarrow$  FV). The validation dataset contains 200 frames.

### The KMeans clustering algorithm

The KMeans clustering algorithm is applied to investigate whether distinct clusters emerge in the dataset. KMeans partitions the data into  $K$  clusters by minimizing the total within-cluster Euclidean distance between data vectors and their respective cluster centroids. This quantity, commonly referred to as the *inertia*, is defined as


$$J = \sum_{k=1}^K \sum_{\underline{x} \in C_k} \left\| \underline{x} - \underline{\mu}_k \right\|^2 \quad (2.1)$$

where  $J$  is the total inertia over all  $K$  clusters,  $C_k$  denotes cluster  $k$ ,  $\underline{x}$  is a data vector belonging to cluster  $C_k$ , and  $\underline{\mu}_k$  is the centroid (mean vector) of cluster  $k$ .

The clustering procedure is initialized with random cluster centroids. The algorithm then iterates between assigning each data vector to the nearest centroid and updating the centroids as the mean of their assigned vectors. This process continues until convergence of the objective function  $J$ .

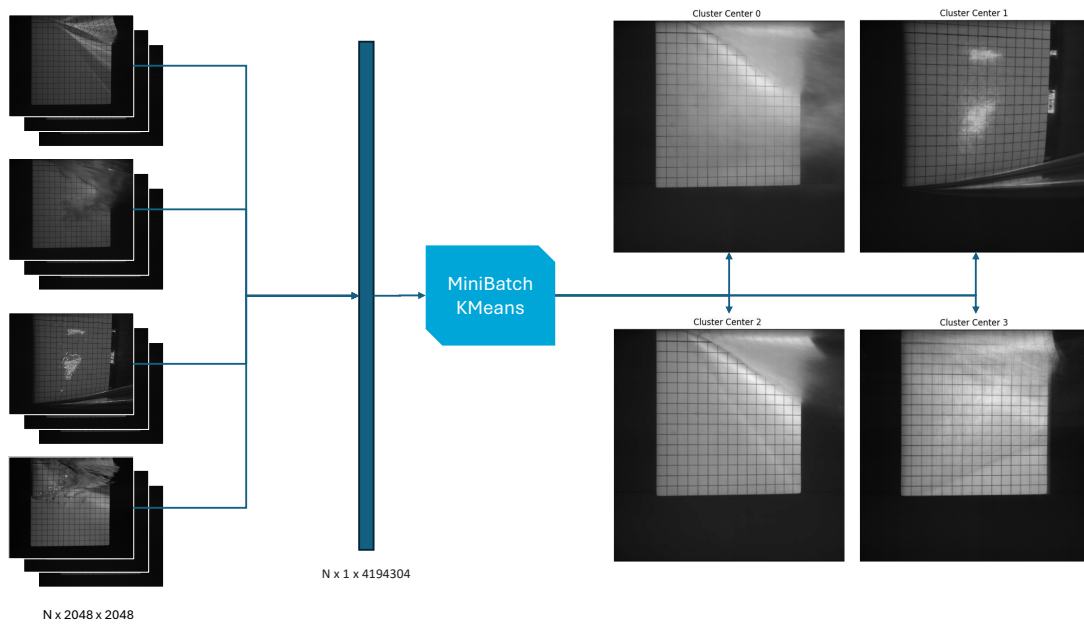
The data vectors are normalized using a standard scaler, meaning they have been transformed such that their mean is zero and their standard deviation is 1.

It is thought that the human interpretation of the following images is at least to some extent captured in the underlying image arrays. The images in Figure 2.7 are typical examples of the respective flow regimes.



**Figure 2.7:** Example frames for different flow regimes

Two strategies were employed to try to find cluster centers representing different flow regimes. Firstly, the naive approach of flattening the image arrays and calculating the cluster centers from the  $N \times 4194304$  vector of grayscale intensities. This approach is graphically explained in Figure 2.8.  $K$  was set to 4, and images were fed from FW, Transitional, PV, and FV.

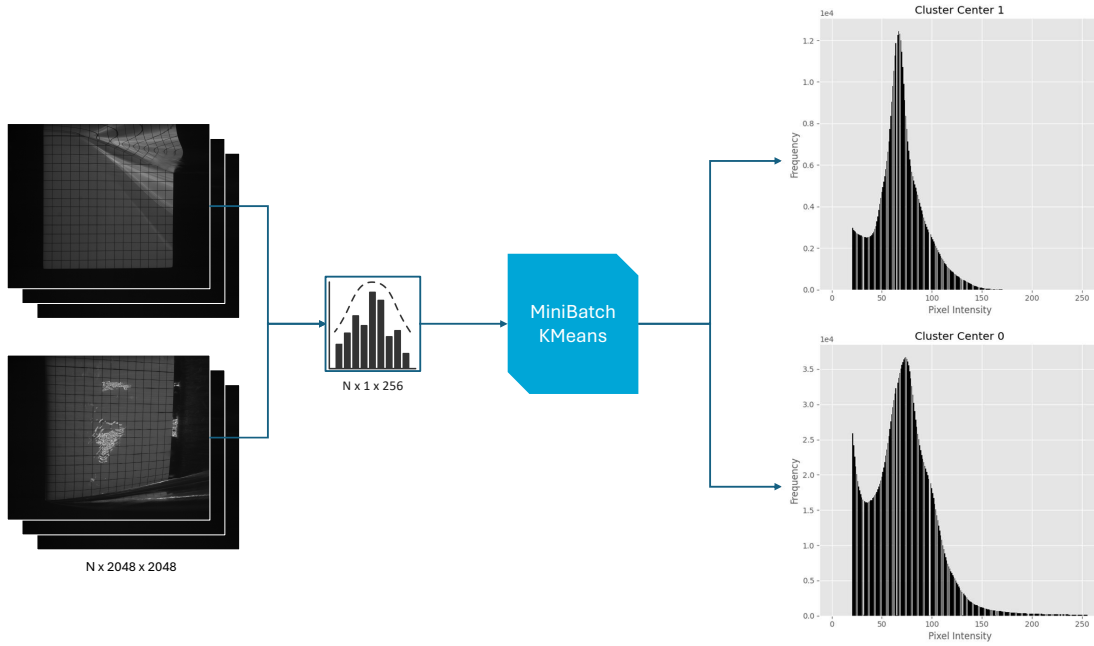


**Figure 2.8:** Clustering based on flattened image arrays with KMeans algorithm  $K = 4$

Secondly, It was hypothesized that the lens glare present in several recordings could introduce a distinct intensity signature in the global gray-level distribution of the images. Rather than treating this optical artifact purely as noise, it was therefore considered as a potentially discriminative feature in the clustering task.

Accordingly, a second clustering strategy was formulated in which the grayscale histogram of each image was computed and used as the feature representation for the KMeans algorithm. A schematic overview of this approach is provided in Figure 2.9.

As an initial proof-of-concept,  $K$  is set to 2, and images are only fed from either FV or FW, providing the largest contrast between the image histograms.



**Figure 2.9:** Clustering based on histogram of the image with KMeans algorithm  $K = 2$

### Performance Metrics and Verification

To assess the quality and physical relevance of the clustering results, multiple quantitative performance metrics are employed. Since the approach is primarily unsupervised, internal validation metrics are used to evaluate cluster compactness and separability. Additionally, where reference labels are available, an external validation measure is introduced. The following section defines each criterion mathematically and provides an intuition for its interpretation.

#### Inertia Contribution

In KMeans clustering, the inertia as defined in (2.1) is the criterion for finding the separation between clusters. The inertia contribution of a single sample  $x_i$  is therefore

$$J_i = \|\underline{x}_i - \underline{\mu}_{C_i}\|^2, \quad (2.2)$$

where  $C_i$  indicates the cluster to which  $\underline{x}_i$  is assigned.

Inertia measures cluster compactness. A low inertia implies that the vectors within a cluster are similar in feature space, which is desirable when seeking physically consistent flow regimes.

#### Silhouette Score

The silhouette score evaluates both intra-cluster cohesion and inter-cluster separation. For a given  $\underline{x}_i$ , define

$$a(i) = \frac{1}{|C_i| - 1} \sum_{\substack{\underline{x}_j \in C_i \\ j \neq i}} \|\underline{x}_i - \underline{x}_j\| \quad (2.3)$$

as the mean distance to all other points in the same cluster, and

$$b(i) = \min_{k \neq i} \left( \frac{1}{|C_k|} \sum_{\underline{x}_j \in C_k} \|\underline{x}_i - \underline{x}_j\| \right) \quad (2.4)$$

as the minimum mean distance to points in another cluster. Note that in (2.3) and (2.4)  $|C_{c(i)}|$  and  $|C_k|$  denote the number of samples of the respective clusters.

The silhouette score of sample  $\underline{x}_i$  is then defined as

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}, \quad (2.5)$$

with  $-1 \leq s(i) \leq 1$ .

A silhouette value close to 1 indicates that a frame is well matched to its assigned regime and clearly distinct from others. Values near 0 suggest a poor distinction between 2 different clusters, and negative values may indicate incorrect assignments [31].

#### Relative Margin to the Next Cluster Centroid

For centroid-based clustering, the distance of a sample  $\underline{x}_i$  to its assigned centroid is

$$d_1(i) = \|\underline{x}_i - \underline{\mu}_{C_i}\|. \quad (2.6)$$

Let  $d_2(i)$  denote the distance to the second closest centroid. The relative margin is defined as

$$r(i) = \frac{d_2(i)}{d_1(i)}. \quad (2.7)$$

The relative margin quantifies classification confidence. If  $r(i) \gg 1$ , the sample lies significantly closer to its own centroid than to the second nearest clusters, indicating a clear regime identification. Values close to 1 imply ambiguity.

#### Confusion Matrix

When partial ground-truth labels are available, the clustering results can be evaluated using a confusion matrix. Let  $y_i$  denote the reference label and  $\hat{y}_i$  the cluster assignment. The confusion matrix  $M$  is defined as

$$M_{pq} = |\{i \mid y_i = p \text{ and } \hat{y}_i = q\}|. \quad (2.8)$$

The confusion matrix reveals how well clusters correspond to known physical regimes. Strong diagonal dominance indicates that clusters align with expected flow categories. Off-diagonal elements may indicate overlapping regimes, transitional behavior, or limitations in the chosen feature representation.

Together, these metrics provide complementary insight: inertia quantifies compactness of a cluster, silhouette score and relative margin measure separability and confidence, and the confusion matrix evaluates physical interpretability when labels are available [5].

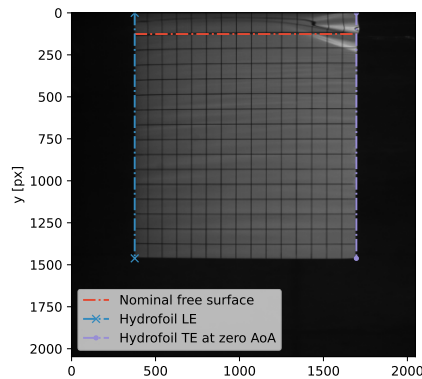
## 2.4. Image Segmentation: Feature Extraction and Tracking

The goal of segmentation is to extract and track meaningful flow structures from each video frame, enabling deeper physical interpretation. For this segmentation task, two main flow features are the objective, namely the free surface of the water for an FW flow, and the cavity closure angle for an FV flow. The image processing techniques used in this algorithm are implemented in the Python library OpenCV [7].

It is important to the Image Segmentation techniques to establish a coordinate system and some geometric constants used in the reconstruction of the free surface and the cavity closure angle.

The image coordinate system has units of pixels and spans 2048 by 2048 pixels. The y-axis is pointing down, and the x-axis is pointing right. The origin is in the top left of the image. The leading edge of the

hydrofoil is at  $x = 376$  px, the nominal free surface when the hydrofoil is at standstill is at  $y = 126$  px, the foil tip leading edge is at  $y = 1462$  px, and the trailing edge of the hydrofoil is at  $x = 1696$  px. These constants are summarized in Figure 2.10




**Figure 2.10:** Constant geometrical features of the NACA hydrofoil for  $AR = 1.0$ , image from run 2004 ( $AoA = 0^\circ$ ,  $F_n = 1.5$ , constant  $AoA$ )

### Free surface drawdown

As mentioned in Section 1.2.1, in an FW flow, the free surface of the water deforms due to wavemaking of the hydrofoil strut. It was deemed relevant to capture and track this deformation of the free surface leading up to the inception of ventilation because it can give an insight into what is observable before the inception of ventilation, and correlate it with what is observed afterwards as the Ventilation Mechanism by Ferreira and Rodrigues [4].

In order to capture the free surface, an algorithm similar to the Hough Gradient Method used by R. Schutt for tracking the horizon, as well as by Stigter et al. for the tracking of micro bubbles, is used [32, 34]. The algorithm was described in Section 1.2.2.

The algorithm consists of 2 pre-processing steps followed by the Probabilistic Hough Transform. Afterwards, filtering and curve fitting is applied. The objective of the algorithm is to locate the grid lines on the hydrofoil as well as their termination points on the free surface. These points can then be used to fit a curve to the free surface. A visual step-by-step diagram of the algorithm is given in Figure 2.11. The mathematical details of each step of the algorithm are further elaborated on in this section.



**Figure 2.11:** Step-by-step algorithm for free surface detection

### Background Subtraction

Background removal is performed using Otsu's method as an initial benchmark. Otsu's method is an unsupervised technique to separate two classes in a grayscale image, namely  $C_0$ , background, and

$C_1$ , foreground. The technique separates the image based on a single pixel threshold  $k$ . Every pixel above the threshold is foreground, and every pixel under the threshold is background. The threshold is determined based on the image histogram. The histogram is hypothesized to be multimodal. Otsu's method determines the threshold that separates these 2 modes based on the between-class variance calculated from the normalized image histogram.

The normalized histogram  $p_i$  of an image with  $L$  grayscale levels (in this case  $L = 256$ ) is defined as (2.9)

$$p_i = \frac{n_i}{N}, \quad \sum_{i=1}^L p_i = 1 \quad (2.9)$$

with  $n_i$  the number of pixels at gray level  $i$  and  $N$  the total number of pixels.

For a candidate threshold  $k$ , the image is partitioned into two classes:

$$C_0 = \{1, \dots, k\}, \quad C_1 = \{k+1, \dots, L\}.$$

The probabilities of being in either class  $C_0$  or  $C_1$  are (2.10).

$$\omega_0(k) = \sum_{i=1}^k p_i, \quad \omega_1(k) = 1 - \omega_0(k), \quad (2.10)$$

and the corresponding class means are (2.11).

$$\mu_0(k) = \frac{1}{\omega_0(k)} \sum_{i=1}^k i p_i, \quad \mu_1(k) = \frac{1}{\omega_1(k)} \sum_{i=k+1}^L i p_i. \quad (2.11)$$

Otsu's method selects the threshold that maximizes the between-class variance  $\sigma_B^2$ , which is a function of the threshold level  $k$ .

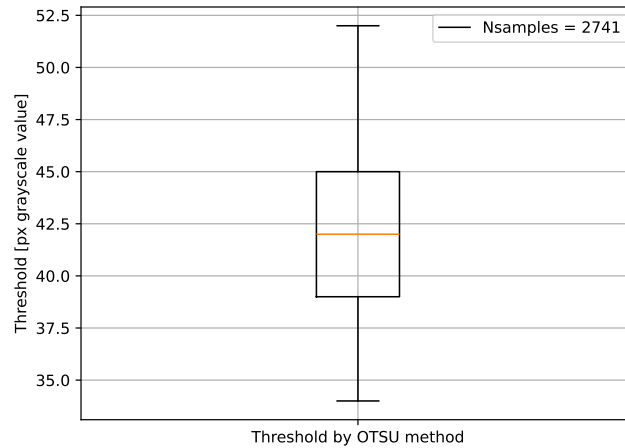
$$\sigma_B^2(k) = \omega_0(k) \omega_1(k) (\mu_1(k) - \mu_0(k))^2.$$

The optimum threshold  $k^*$  is therefore found by evaluating.

$$k^* = \arg \max_{1 \leq k < L} \sigma_B^2(k).$$

The full derivation and existence of solutions is elaborated in [29].


To assess the robustness of this approach, Otsu's threshold was computed for all frames labeled as fully wetted (FW), yielding a total of 2741 samples. The resulting distribution of the optimal threshold for each sample is shown in Figure 2.12.



**Figure 2.12:** Optimal threshold from Otsu's method for background removal in FW images (total 2741 frames)

The boxplot demonstrates that the automatically selected threshold values are tightly distributed around a mean grayscale value of approximately 42, with limited variance. This indicates that the illumination conditions and background intensity levels are consistent throughout the dataset. The narrow spread of threshold values suggests that a static threshold is feasible and would yield stable segmentation results.

Although the minimum Otsu threshold observed is approximately 34, visual inspection reveals that such a threshold may remove physically relevant shadow regions near the leading edge and in the wake of the hydrofoil. Figure 2.13 compares background masks generated with thresholds of 20 and 34. A threshold of 34 removes a larger portion of low-intensity regions, including parts of the leading edge structure, whereas a threshold of 20 better preserves these features while still effectively suppressing the background.



**Figure 2.13:** Examples of a frame from run 2080 ( $F_n = 2.5$ ,  $AR = 1.0$ , variable AoA) frames and masks with different background removal thresholds

Since the objective of pre-processing is to remove the static background without affecting flow-related structures, a conservative static threshold of 20 is adopted for the full pre-processing pipeline. This value ensures robust background suppression while minimizing interference with physically meaningful image content. An additional benefit is that the static threshold is computationally cheaper and will therefore yield a faster algorithm.

### Gaussian blurring

A Gaussian blurring kernel is applied to reduce high-frequency noise. It was found that a 7x7 kernel with coefficients from the standard normal distribution strikes the right balance between preserving edges and reducing noise.

A 7x7 Gaussian kernel is given in (2.12)

$$G = \frac{1}{273} \begin{bmatrix} 1 & 4 & 7 & 10 & 7 & 4 & 1 \\ 4 & 16 & 26 & 33 & 26 & 16 & 4 \\ 7 & 26 & 55 & 71 & 55 & 26 & 7 \\ 10 & 33 & 71 & 91 & 71 & 33 & 10 \\ 7 & 26 & 55 & 71 & 55 & 26 & 7 \\ 4 & 16 & 26 & 33 & 26 & 16 & 4 \\ 1 & 4 & 7 & 10 & 7 & 4 & 1 \end{bmatrix} \quad (2.12)$$

The blurred image is given by the convolution of the original image (with the background mask) with the kernel (2.13).

$$I' = G * I \quad (2.13)$$

### Canny Edge Detection

The Canny edge detection algorithm is used to extract edge information from the input images. The input to the Canny algorithm is a grayscale image with the background removed and Gaussian blurring applied to reduce noise.

Edge detection is then performed by evaluating intensity gradients. The horizontal and vertical gradients are computed using the Sobel operators  $S_x$  and  $S_y$ , which provide discrete approximations of first-order derivatives while incorporating local smoothing. The Sobel kernels can be interpreted as a central difference operator in one direction combined with weighted averaging in the orthogonal direction. This structure improves robustness to noise compared to only a 1x3 finite difference scheme. The Sobel kernels in the  $x$  and  $y$  directions are given by

$$S_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}, \quad S_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}. \quad (2.14)$$

The gradient components are computed by convolution,

$$G_x = S_x * I, \quad G_y = S_y * I. \quad (2.15)$$

From these components, the gradient magnitude and orientation are obtained as

$$|\nabla I| = \sqrt{G_x^2 + G_y^2}, \quad \theta = \arctan\left(\frac{G_y}{G_x}\right). \quad (2.16)$$

The gradient magnitude represents the strength of the local intensity change and forms the basis for edge detection. Non-maximum suppression is applied along the gradient direction to retain only local maxima. This process ensures that thick and blurry regions identified by the Sobel operators are thinned down to a single-pixel width by suppressing all pixels that do not constitute a local peak in gradient intensity.


The algorithm is controlled primarily by two threshold parameters, a lower threshold *low\_threshold* and an upper threshold *high\_threshold*, which determine how strong an intensity gradient must be to be classified as an edge. Pixels with  $|\nabla I| > \text{high\_threshold}$  are immediately classified as strong edges. Pixels with  $\text{low\_threshold} < |\nabla I| \leq \text{high\_threshold}$  are classified as weak edges and are retained only if they are connected to a strong edge through edge continuity. This process is called hysteresis thresholding. Pixels with  $|\nabla I| \leq \text{low\_threshold}$  are suppressed.

Note that the theoretical maximum value of the Sobel operator  $S_x$  or  $S_y$  is 1020 (on an 8-bit image). This means the theoretical maximum of  $|\nabla I|$  is 1442. In real images, however, the intensity gradient

is significantly lower because an image never has perfect contrast vertical or horizontal gradients.

The output of the Canny algorithm is a binary edge map in which edge pixels are given a value of one and all other pixels are set to zero. This edge map provides a clean and well-defined representation of linear features, making it a suitable and reliable input for subsequent line segment detection using the Probabilistic Hough Transform.

For correct detection of the grid lines, a lower threshold of 30 and an upper threshold of 100 were used. An example output of the edge detection is given in Figure 2.14.



**Figure 2.14:** Example of canny edge detection applied to frame nr 2082 (T - 0.5 s) of run 2044 ( $F_n = 1.5$ , AR = 1.0, variable AoA).

The majority of relevant edges are successfully captured, providing sufficient information for reconstruction of the free surface. The subsequent step is to convert this binary edge image into coordinate-based line information using the Probabilistic Hough Transform.

#### Probabilistic Hough Transform

The Probabilistic Hough Transform is used for line segment detection due to its robustness and computational efficiency when identifying finite, well-defined line features. Unlike the standard Hough Transform, which detects infinite lines, the Probabilistic Hough Transform directly returns line segments with defined start and end points, making it better suited for locating the grid lines and their termination at the free surface. On the hydrofoil grid, only partial line segments are visible, and the endpoints are required for the free surface reconstruction. The algorithm takes in the edge view from Canny Edge Detection and translates it into an array of line segments stored endpoint-to-endpoint.

Mathematically, the Hough Transform is based on the parametric representation of a line in polar coordinates,

$$\rho = x \cos \alpha + y \sin \alpha \quad \text{or} \quad y = \left( -\frac{\cos \theta}{\sin \theta} \right) x + \left( \frac{\rho}{\sin \theta} \right) \quad (2.17)$$

where  $\rho$  is the perpendicular distance from the origin to the line and  $\alpha$  is the angle between the normal vector of the line and the  $x$ -axis. Each edge pixel  $(x, y)$  in the binary edge image defines a sinusoidal curve in the  $(\rho, \alpha)$  parameter space. Intersections of these curves correspond to parameters  $(\rho, \alpha)$  that are supported by multiple edge pixels, indicating the presence of a line.

In the Probabilistic Hough Transform, a random subset of edge pixels is sampled instead of using all edge points. For each sampled edge pixel, candidate lines are evaluated in the discretized  $(\rho, \alpha)$  space. A line hypothesis is accepted if the number of supporting edge pixels exceeds a predefined threshold. Once a valid line is detected, connected edge points along this direction are traced to determine the start and end coordinates of the corresponding line segment. This results in a finite line

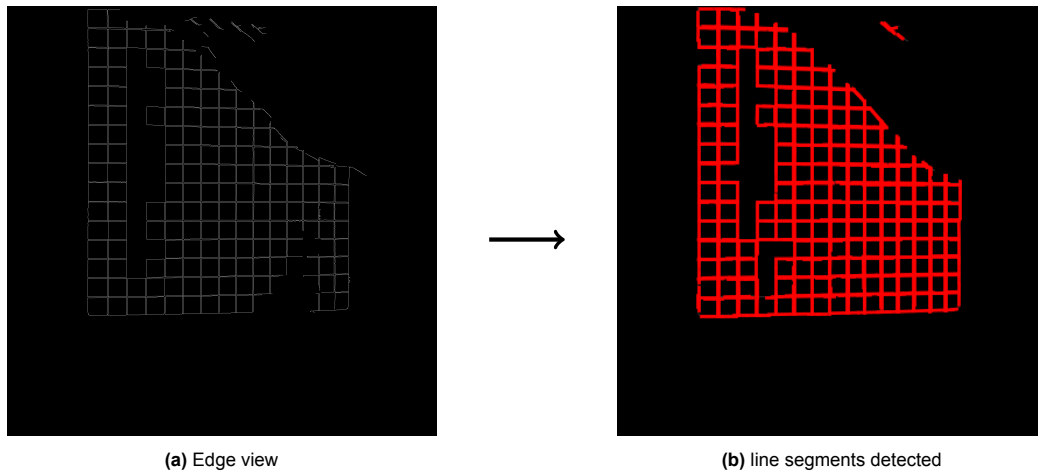
segment defined by two endpoints rather than an infinite line. The exact details of the implementation are found in [26].

The Probabilistic Hough Transform depends on several parameters that control how line segments are detected:  $\rho$ ,  $\alpha$ , threshold, minLineLength, and maxLineGap. The parameters  $\rho$  and  $\alpha$  define the discretization resolution of the parameter space in terms of distance and angle. The threshold sets the minimum number of edge points required before a line hypothesis is accepted. The parameter minLineLength specifies the shortest line segment that is considered valid, helping to discard short, spurious detections, while maxLineGap defines the largest gap allowed between collinear edge points for them to be treated as a single continuous line segment.

These parameters were tuned empirically by experimenting on representative images. A set of parameters was found that yields good and consistent results for the present application; however, it cannot be concluded that this set is optimal. The values used are summarized in Table 2.3. An example of the algorithm's output can be viewed in Figure 2.15.

**Table 2.3:** Parameter values for Probabilistic Hough Transform used in Free surface tracking

<i>Parameter</i>	<i>Value</i>
$\rho$	1
$\alpha$	$\pi / 180$ [rad]
Threshold	10 [-]
minLineLength	20 [px]
maxLineGap	15 [px]



**Figure 2.15:** Example of Probabilistic Hough Transform applied to frame nr 2082 (T - 0.5 s) of run 2044 ( $F_n = 1.5$ , AR = 1.0, variable AoA).

It can be seen that the Hough Gradient method correctly captures the grid lines in the edge view of the image. However, because the grid drawn on the hydrofoil is multiple pixels wide, the canny edge detection actually produces two lines of high intensity gradient. It is thus concluded that additional filtering is needed to clean up the grid detection.

#### Filtering and combining of houghlines

The Hough gradient method, as described thus far, yields line segments stored endpoint to endpoint from the edge view. Note that because the Hough Gradient Method failed to connect all line segments that are on the same line, it was decided to use a filtering algorithm to connect the dots anyway. The filter is threefold. Firstly, the Hough lines are sorted in 3 groups, namely Horizontal (all Hough lines in the angle band  $-5$  to  $5$  [°]), Vertical (all Hough lines in the angle band  $85$  to  $95$  [°]), and Oblique (All remaining line segments). The oblique lines are discarded. The vertical lines are compared by their

---

x coordinate. If two Hough line segments have an x coordinate within a margin of 5 pixels, they are flagged as potentially on the same vertical line. To further illustrate how exactly the Houghlines are combined and filtered, the pseudocode in Algorithm 1 explains it exactly.

**Algorithm 1** Classify and Merge Hough Line Segments

---

**Require:** HoughLines =  $\{(x_1, y_1, x_2, y_2)\}$ ,  $TE\_X\_MAX$   
**Ensure:** CombinedVerticalLines, CombinedHorizontalLines, ObliqueLines, MiscellaneousLines


- 1: Initialize empty sets:
- 2: HorizontalLines, VerticalLines, ObliqueLines, MiscellaneousLines
- 3: % \_\_\_\_\_
- 4: % STEP 1: CLASSIFICATION
- 5: % \_\_\_\_\_
- 6: **for** each line  $l = (x_1, y_1, x_2, y_2)$  in HoughLines **do**
- 7:   Compute orientation angle  $\theta$
- 8:   **if**  $x_1 > TE\_X\_MAX$  OR  $x_2 > TE\_X\_MAX$  **then**
- 9:     Add  $l$  to MiscellaneousLines
- 10:   **end if**
- 11:   **if**  $|\theta| < 5^\circ$  **then**
- 12:     Add  $l$  to HorizontalLines
- 13:   **else if**  $|\theta| > 85^\circ$  **then**
- 14:     Add  $l$  to VerticalLines
- 15:   **else**
- 16:     Add  $l$  to ObliqueLines
- 17:   **end if**
- 18: **end for**
- 19: % \_\_\_\_\_
- 20: % STEP 2: PROCESS VERTICAL LINES
- 21: % \_\_\_\_\_
- 22: **for** each line in VerticalLines **do**
- 23:   % Standardize endpoint order
- 24:   Ensure  $y_1 \leq y_2$
- 25: **end for**
- 26: Sort VerticalLines by  $x$ -coordinate
- 27: Group neighbouring lines with  $|\Delta x| < 5$
- 28: **for** each group  $G$  **do**
- 29:    $x \leftarrow$  mean  $x$  of  $G$
- 30:    $y_{min} \leftarrow$  minimum lower endpoint in  $G$
- 31:    $y_{max} \leftarrow$  maximum upper endpoint in  $G$
- 32:   Add merged line  $(x, y_{min}, x, y_{max})$  to CombinedVerticalLines
- 33: **end for**
- 34: % \_\_\_\_\_
- 35: % STEP 3: PROCESS HORIZONTAL LINES
- 36: % \_\_\_\_\_
- 37: **for** each line in HorizontalLines **do**
- 38:   % Standardize endpoint order
- 39:   Ensure  $x_1 \leq x_2$
- 40: **end for**
- 41: Sort HorizontalLines by  $y$ -coordinate
- 42: Group neighbouring lines with  $|\Delta y| < 5$
- 43: **for** each group  $G$  **do**
- 44:    $y \leftarrow$  mean  $y$  of  $G$
- 45:    $x_{min} \leftarrow$  minimum left endpoint in  $G$
- 46:    $x_{max} \leftarrow$  maximum right endpoint in  $G$
- 47:   Add merged line  $(x_{min}, y, x_{max}, y)$  to CombinedHorizontalLines
- 48: **end for**
- 49: **return** CombinedVerticalLines, CombinedHorizontalLines, ObliqueLines, MiscellaneousLines

---

It was decided to only use the vertical grid lines because they are not subject to the perspective shift of the variable angle of attack. In addition, they do not show a curvature from the hydrofoil cross-section.

The vertical grid lines thus result in more reliable detection and more accurate endpoint coordinates.

The combined and filtered houghlines can be seen in Figure 2.16.



**Figure 2.16:** Example of combined and filtered houghlines from frame nr. 2082 (T - 0.5 s) of run 2044


It can be seen that the vertical grid lines are correctly reconstructed but the horizontal line segments suffer from the perspective shift of the hydrofoil curvature and thus result in multiple hough line segments detected for one an the same grid line on the hydrofoil.

#### Free surface points and curve fitting

The shape of the free surface is well captured by an asymmetric sigmoid function (2.18). The ordinary least squares algorithm was used to find a set of parameters  $\{A, k, x_0, C\}$  that minimizes the residuals between the curve and the detected free surface points. The detected points were first scaled from image coordinates to normalized coordinates using a MinMax scaling. The scaling improves the robustness of the algorithm.

$$\text{sgf}(x) = \frac{A}{1 + e^{k(x-x_0)}} - C \quad (2.18)$$

An example of the curve fitting is given in Figure 2.17



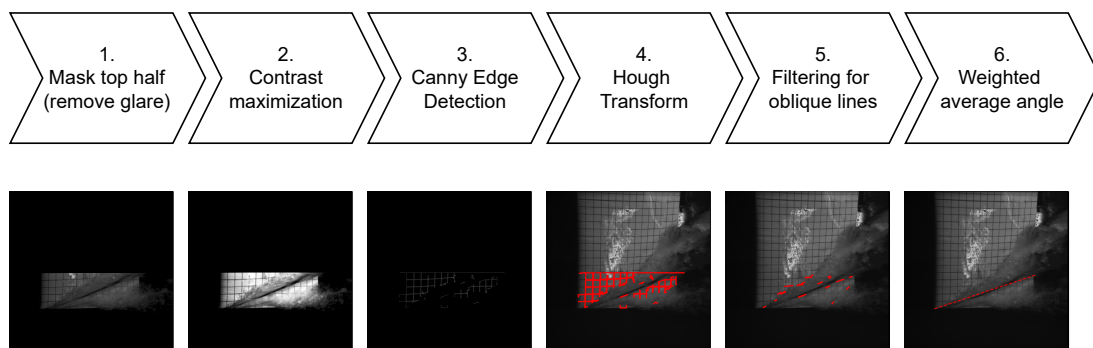
**Figure 2.17:** Example of curve fitting of the detected points to frame nr 2082 (T - 0.5 s) of run 2044. Curve parameters found:  $A = 1.124$ ,  $k = -6.170$ ,  $x_0 = 0.595$ ,  $C = 0.024$ . RMSE = 0.017

### Cavity closure angle reconstruction

As discussed in Section 1.2.1, the cavity closure angle ( $\Phi$ ) is a key geometric feature associated with a fully ventilated flow. Quantifying this angle provides insight into the cavity structure near the hydrofoil tip and enables a consistent comparison between different flow conditions. Accurate detection of the cavity boundary is therefore required in the vicinity of the foil tip. An additional difficulty arises due to the air to water interface which introduces glare.

To determine the cavity closure angle, an image processing algorithm similar to the free surface tracking approach is employed. The primary difference lies in the definition of the region of interest, which is shifted towards the hydrofoil tip to avoid contrast spikes caused by glare. Within this region, additional contrast enhancement is applied to obtain a sharper definition of the cavity boundary prior to edge detection.

The algorithm then follows the same general structure as the free surface detection method. Canny edge detection is applied to extract edge information, after which the Probabilistic Hough Transform is used to identify line segments corresponding to the cavity interface. Finally, a set of filtering operations is performed to retain only the physically relevant line segments, from which the cavity closure angle is calculated. A schematic overview of the algorithm is provided in Figure 2.18.




**Figure 2.18:** Step-by-step algorithm for cavity closure angle detection

#### glare removal and contrast maximization

As can be seen in various images in this report, for example Figure 2.4f. An image artifact appears that has a profound effect on the usability of image segmentation techniques. The sharp spikes in contrast due to lens glare from the air to water interface are a major source of spurious detections. Because the glare is most pronounced in FV flows, it was decided to tackle it pragmatically, by shifting the region of interest to the foil tip.

Figure 2.19 shows the steps taken to get a clean image in which the cavity closure angle can be reliably reconstructed. The contrast maximization is performed by histogram equalization.




**Figure 2.19:** Example contrast maximization run 2044 ( $F_n = 1.5$ , AR = 1.0, variable AoA) at  $T + 0.4$  s (frame nr 2172)

This workflow was motivated by the presence of glare, which introduces regions of high-intensity pixels in the image. These saturated regions dominate the intensity range and reduce the effectiveness of the Canny edge detection thresholds. As a result, the contrast associated with the cavity boundary becomes relatively weak and is not reliably detected. Without glare removal and contrast enhancement, the cavity closure edge cannot be consistently identified.

#### Edge detection and Hough transform

Similar to the Hough Gradient Method used in the free surface reconstruction, the method here attempts to find an edge related to the flow. In this case the cavity closure angle. An example detection is shown in Figure 2.20. Note that no blurring was applied because all the contrast is necessary to obtain a detection. Note also the long, consistent Hough line segment at  $y = 1024$  px due to the sharp contrast by slicing the top of the image off. A background removal threshold at 20 px was still applied, similar to the free surface reconstruction.



**Figure 2.20:** Example of Probabilistic Hough Transform applied to frame 2172 of run nr 2044


The parameters for the Canny edge detection and Probabilistic Hough Transform for the cavity closure angle reconstruction are given in Table 2.4

**Table 2.4:** Parameter values for Probabilistic Hough Transform used in cavity closure angle

<i>Parameter</i>	<i>Value</i>
canny low_threshold	0
canny high_threshold	255
$\rho$	1
$\alpha$	$\pi / 180$ [rad]
Hough threshold	15 [px]
minLineLength	20 [px]
maxLineGap	15 [px]


#### Filtering for oblique line segments

As shown in Figure 2.20, the grid is once again captured by the Probabilistic Hough Transform. A bandpass filter with a band between  $5^\circ$  and  $45^\circ$  is applied to the detected Hough line segments. As mentioned in Section 1.2.1, the maximum cavity closure angle for a stable ventilated cavity is  $45^\circ$ . The lower bound of  $5^\circ$  was found to be the smallest value that clearly distinguishes between a horizontal grid line and an oblique line. The result is all detected Hough lines at an oblique angle. The result of the filter is displayed in Figure 2.21

**Figure 2.21:** Example of bandpass filter [ $5^\circ$ ,  $45^\circ$ ] applied to frame 2172 of run nr 2044

#### Weighted average angle

To derive a single cavity closure angle from the set of detected oblique Hough line segments, a weighted averaging approach is employed. To be precise, the weighted average angle of all line segments flagged as oblique, with the weight being their length, is calculated. This approach is based on the assumption that the cavity boundary forms a long, continuous edge, whereas the spray and spurious features are typically represented by shorter line segments. Weighting by line length, therefore, emphasizes the physically relevant cavity interface while reducing the influence of short, noisy detections. The resulting cavity closure angle is shown in Figure 2.22.



**Figure 2.22:** Detected cavity closure angle of frame 2172 from run 2044.  $\Phi = 18.47^\circ$

# 3

## Results

This chapter presents the results of the developed methodologies in three main parts. In Section 3.1, the performance of the flow regime clustering is evaluated using KMeans clustering, comparing the full-image feature representation with the histogram approach.

Second, in Section 3.2, the reconstruction of the free surface is presented, obtained through application of the Hough gradient method to the image data.

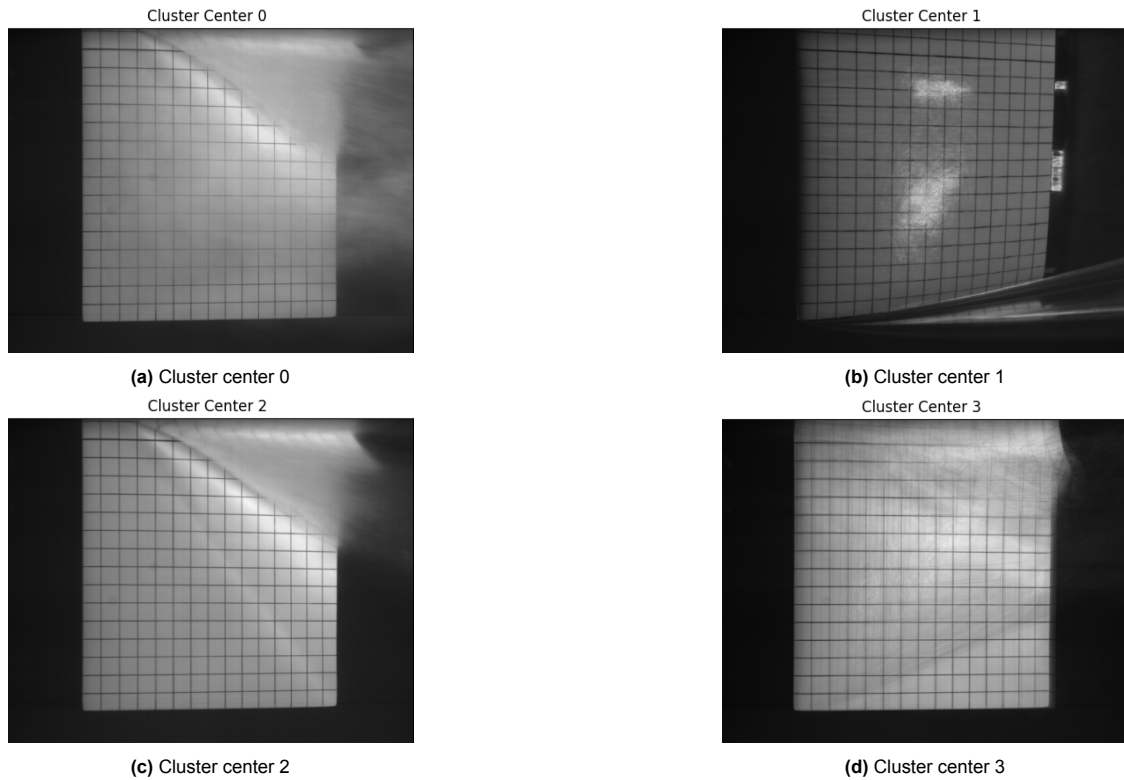
Finally, in Section 3.3, the extraction of the cavity closure angle is discussed. This geometric quantity serves as a physically interpretable parameter derived from the reconstructed air-to-water interface and provides further insight into the ventilated cavity.

### 3.1. Flow Regime Classifier

The results from the clustering algorithms as proposed in Section 2.3 are presented here. Specifically, the performance of each algorithm is evaluated based on its ability to accurately distinguish between different flow regimes and maintain consistency across the image dataset.

#### Flattened Image Clustering

These cluster centers are the result of the initial proposal from Section 2.3. Namely, to construct 4 clusters from a balanced stratified random sample of images belonging to the flow regimes FV, PV, FW, or Transitional. They are the same centers as presented in the diagram Figure 2.8 from Section 2.3. The cluster centers are represented in Figure 3.1.




**Figure 3.1:** Cluster centers from flattened image clustering  $k = 4$

Visual inspection of the cluster centroids suggests partial correspondence with physically meaningful flow regimes. In particular, clusters 0, 1, and 2 exhibit visual characteristics that resemble transitional flow, fully ventilated (FV) flow, and fully wetted (FW) flow, respectively. However, cluster 3 does not display features consistent with partial ventilation (PV), indicating that not all regimes are adequately represented.

Despite the apparent qualitative agreement for several clusters, quantitative evaluation using the metrics defined in (2.2), (2.5), and (2.7) indicates that the clustering does not reliably distinguish between the supplied flow regimes.


Firstly, A scatter plot of silhouette score vs the relative margin score is presented in Figure 3.2. The objective of this plot is to identify if the cluster centers are well separated in feature space. As discussed in Section 2.3, the silhouette score ranges from -1 to 1, and the relative margin ranges from 0 to all positive values. Note that the top right corner in this plot indicates samples that have been assigned to a cluster with high confidence.



**Figure 3.2:** Silhouette score vs relative margin for the  $k = 4$  flattened image model, scores obtained from the samples in the validation set

It can be seen that the relative margin is never higher than 1, indicating poor separation between the clusters. Also note that the silhouette score is on average 0.2, which is also insufficient.

To further analyse the clustering, a histogram of the inertia contributions of each sample is evaluated. The inertia contribution indicates whether the samples are close in feature space to their assigned cluster centroid. It is presented in Figure 3.3



**Figure 3.3:** Histogram of the inertia contribution for the  $k = 4$  flattened image model, obtained from the samples in the validation set

The distribution has a clear peak at the lower contributions and relatively few higher-order contributions. This indicates that the majority of validation samples are geometrically close to their assigned cluster centers, suggesting reasonably compact clusters.


The evaluation of the silhouette score and relative margin, combined with the inertia contributions, indicates limited cluster separability. A substantial portion of the samples exhibits low silhouette values and relative margins, suggesting ambiguity in cluster assignment and overlap in feature space. This lack of clear geometric separation implies that the clustering does not provide a robust partitioning of

the data into distinct flow regimes.

Furthermore, while cluster centroids 0, 1, and 2 exhibit visual resemblance to physically interpretable regimes, cluster 3 cannot be consistently associated with a known ventilation state. The absence of a clear physical interpretation for this cluster complicates comparison with human labels and undermines the practical applicability of the model.

It is therefore concluded that the flattened-image KMeans approach with  $k = 4$  is not sufficiently reliable for automated classification of flow regimes.

In addition to the limited separability observed in the validation metrics, this raises the question of whether the chosen number of clusters,  $k = 4$ , is appropriate for the intrinsic structure of the data. To investigate this, the variation of the total inertia with respect to the number of clusters is analysed.



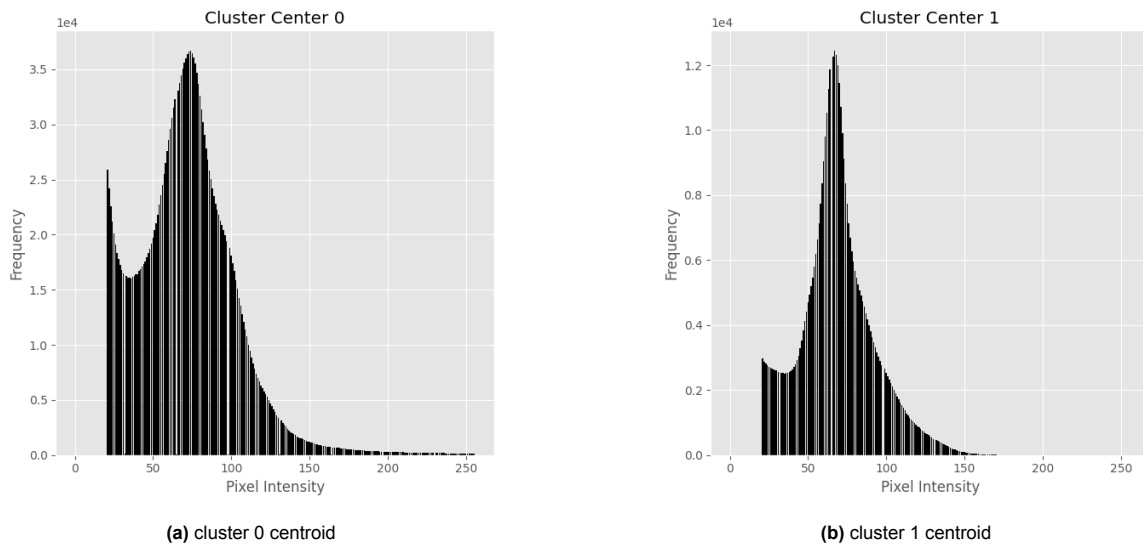
**Figure 3.4:** Variation of number of clusters  $k$  and the resulting Inertia

Figure 3.4 reveals a pronounced change in slope at  $K = 3$ . For  $K > 3$ , the decrease in inertia becomes progressively smaller, indicating diminishing returns from additional clusters. Therefore,  $K = 3$  is proposed as the optimal number of clusters for a reiteration of this clustering.

#### Histogram-based Clustering

As described in Section 2.3, a histogram-based approach was applied as a potential way to identify the high intensity pixels in a Fully Ventilated flow with lens glare, and thereby distinguish it from the other flow regimes.

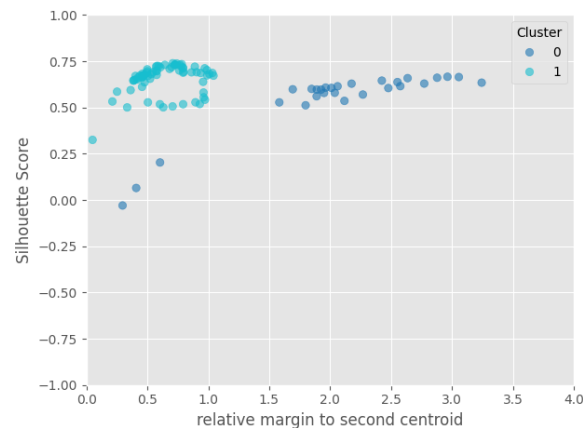
The cluster centers of a  $k = 2$  clustering are shown in Figure 3.5. They are the same centers as presented in the diagram Figure 2.9.



**Figure 3.5:** Cluster centers from KMeans histogram clustering  $k = 2$

It can be seen that cluster 0 has more high-intensity pixels, especially more pixels above a grayscale value of 150. It is therefore concluded that cluster 0 maps to the human label of 'FV'. For the rest, the distributions look similar in shape.


Again, the silhouette score vs relative margin is evaluated to identify the separation between the cluster and the reliability of the assignment. The result is shown in Figure 3.6.



**Figure 3.6:** Silhouette score vs relative margin for the  $k = 2$  histogram model, scores obtained from the samples in the validation set

Note that cluster 0 has a higher margin on average than cluster 1. This is expected because the training data was sampled to maximize the difference between 'FV' and 'Not FV'. This figure indicates that there is a clear decision boundary between the two groups 'FV' and 'Not FV'

Figure 3.7 shows the inertia contributions of the evaluated samples. It can be seen that the distribution is not as cleanly grouped to the left.




**Figure 3.7:** Histogram of the inertia contribution for the  $k = 2$  histogram model, obtained from the samples in the validation set

Although the training dataset was constructed to contain an equal number of fully ventilated and non-fully ventilated frames, the inertia distribution of the  $K = 2$  histogram-based model exhibits a broad spread with substantially high-contribution samples. This behaviour can be explained by the internal heterogeneity of the non-FV class, which comprises fully wetted, partially ventilated, and transitional frames.

By enforcing a single centroid to represent this structurally diverse group, the model necessarily produces large intra-cluster variance. The resulting centroid represents an average intensity distribution that does not accurately characterize any single physical regime, leading to elevated inertia contributions for many samples.

This indicates that the  $k = 2$  histogram-based representation is insufficient to capture the underlying physical variability of the flow. This is as expected. The original hypothesis was that the clustering could exploit the lens glare as an image artifact to distinguish the flow regime.

Finally, because this clustering is capable of distinguishing a human-interpretable flow regime from the rest, a confusion matrix is constructed. The confusion matrix is shown in Figure 3.8.



**Figure 3.8:** Confusion Matrix obtained from the validation data using the  $k = 2$  histogram clustering.

It can be seen that the number of true negatives and the number of true positives are acceptable. Especially considering that there were 20 FV frames in the validation dataset, out of a total of 200 frames. The number of false positives is very high. Likely because a partial cavity also exhibits lens glare.


## 3.2. Free Surface Tracking

Section 2.4 explained the methodologies employed to reconstruct the free surface of the flow and the cavity closure angle from the image dataset.

The free surface reconstruction method has been applied to every frame labeled as fully wetted in the image dataset. This is a total of 4021 images over 64 runs.

For each image, the sigmoid parameters and Root Mean Square Error from the curve fit were recorded. There was also a record of the free surface drawdown at the trailing edge prior to ventilation in the data by Ferreira and Rodrigues. This was deemed as an appropriate validation criterion because the reduced dataset only contains images just prior to ventilation, as explained in Section 2.2

The error is defined as the difference between the free surface drawdown at the trailing edge obtained from the fitted curve and the manually identified drawdown prior to ventilation, as determined independently by Ferreira and Rodrigues. The distribution of this error and the corresponding RMSE values are shown in Figure 3.9.



**Figure 3.9:** Histogram of the error and RMSE of the free surface curve fitting

The horizontal axis is expressed both in pixels and in hydrofoil grid cell units, where one grid cell corresponds to approximately 80 px. Most realizations exhibit an absolute error within approximately two grid cells, while the majority of RMSE values are below 0.1. This indicates that, for most frames, the fitted free surface profile reproduces the manually observed drawdown within a limited spatial deviation.

Based on the observed RMSE distribution, a threshold of  $RMSE < 0.1$  was adopted as a selection criterion for inclusion in the ensemble averaging procedure. This threshold excludes fits associated with visibly poor convergence or incomplete free surface detection, while retaining the majority of physically consistent realizations.

Figure 3.9 demonstrates that the free surface procedure described in Section 2.4 provides reproducible

estimates of the trailing-edge drawdown and that the RMSE serves as a practical metric for assessing detection reliability.

The number of frames affected by the applied threshold, together with the number of frames resulting in an ill-conditioned fit, is summarized in Section 3.2. An ill-conditioned fit refers to either of the following situations: (i) the Probabilistic Hough transform fails to detect any candidate points, or (ii) the detected points lead to a non-convergent or numerically unstable fit of the sigmoid function. In both cases, the free surface reconstruction procedure is unsuccessful.

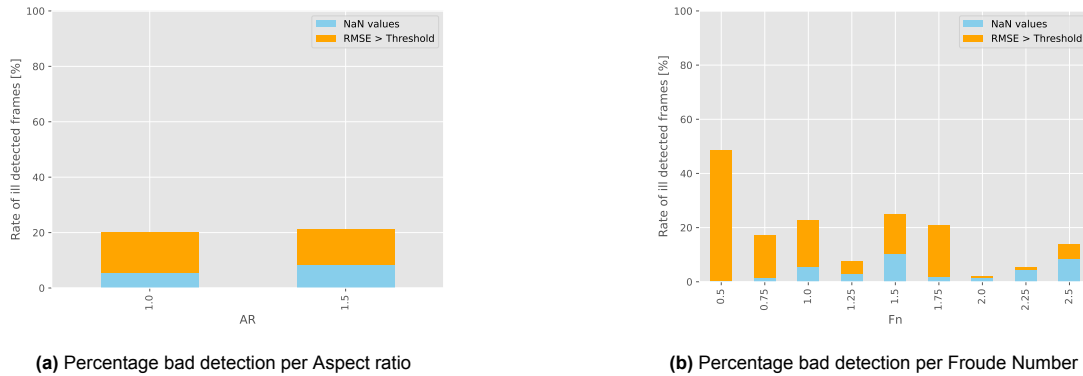


Figure 3.10: Summary of the success percentage of the Free Surface Reconstruction.

It can be seen that, on average, the free surface reconstruction succeeds on 80% of the frames. Furthermore, there is no observable difference in effectiveness between the two geometries, namely, an Aspect ratio of 1 and an aspect ratio of 1.5. Finally, it is interesting to note that the number of ill-detected frames is higher at low Froude numbers. It is hypothesized that this is due to the very small free surface drawdown associated with this Froude Number regime. The sigmoid function cannot reliably follow such a contour.

To study the relationship between the free surface shape and the ventilation onset, a plot is made of the average free surface sigmoid for the different ventilation mechanisms identified by Ferreira and Rodrigues. The result is shown in Figure 3.11.


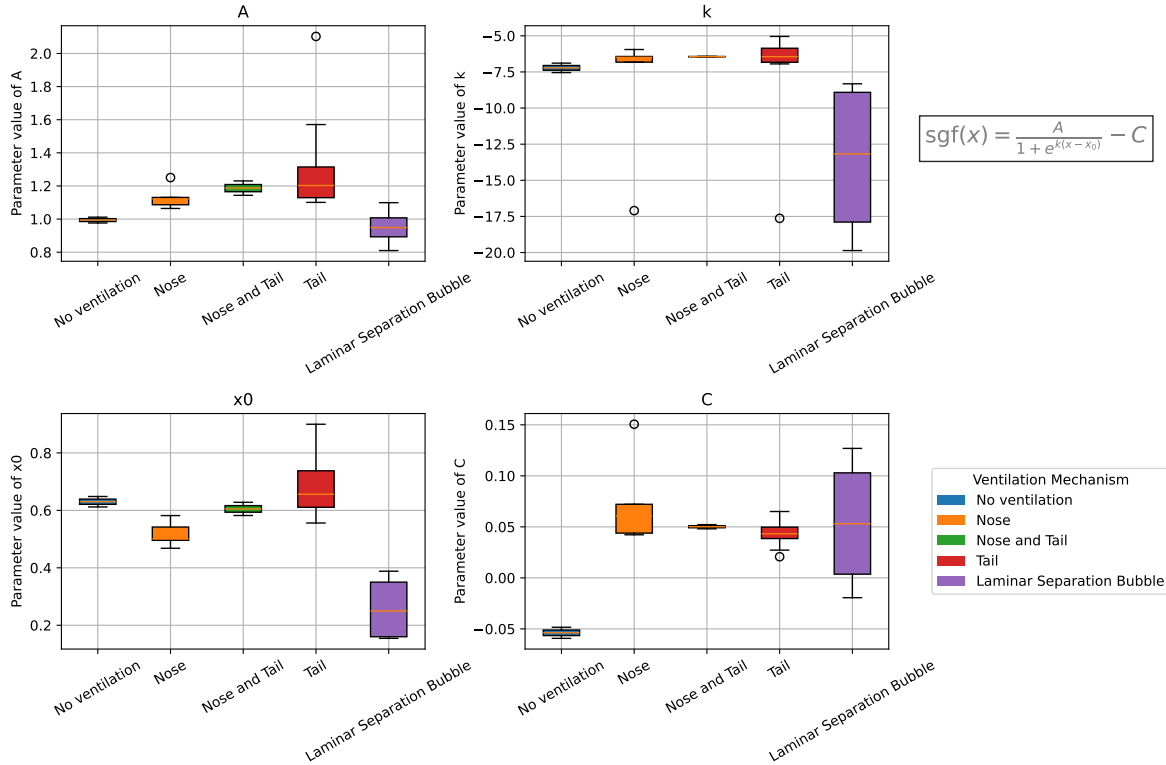


Figure 3.11: Per Ventilation mechanism mean asymmetric sigmoid curves for every frame before ventilation

It can be seen that there is no meaningful difference between the cases 'No Ventilation', 'Nose' Ventilation, 'Nose and Tail' ventilation, and 'Tail' ventilation, but there is a pronounced difference in free surface shape for Ventilation due to a 'Laminar Separation Bubble'.

To further investigate this finding, a plot is presented showing the mean and variance of each parameter  $a$ ,  $k$ ,  $x_0$ ,  $C$  in the sigmoid function for the different ventilation mechanisms. It can be found in Figure 3.12



**Figure 3.12:** Per Ventilation mechanism mean of the parameters in the asymmetric sigmoid curve for every frame before ventilation


This figure confirms that there is no significant observable difference between the ventilation mechanisms except 'Laminar Separation Bubble'. This is remarkable because the shape of the free surface is a result of the pressure distribution on the hydrofoil and the LSB has a significant effect on the pressure distribution. Note that from Figure 3.12 it can be seen that  $x_0$  is significantly smaller for the LSB group. Also note that  $k$  is significantly more negative for the LSB group compared to the other mechanisms. Physically a larger  $x_0$  and a more negative  $k$  mean that the free surface depresses further towards the leading edge and more sharply in the depth direction compared to the other cases. Concerning here is the large variance in  $k$  and  $C$  for the LSB group.

### 3.3. Cavity Closure Angle

Finally, the algorithm for the reconstruction of the cavity closure angle is evaluated. The objective of the methodology is to reconstruct the cavity closure angle from the sharp contrast at the air to water interface. There is no validation data for this dataset, thus a physics based verification method is applied.

The physics of the flow as discussed in Section 1.2.1 are such that at higher speeds (higher Froude numbers), the cavity closure angle is smaller. This is as a consequence of the re-entrant jet, where higher Froude numbers are associated with more elongated cavity structures and consequently smaller closure angles. Therefore, a plot of the detected cavity closure angle for all samples labeled as 'FV'

was applied. The total number of samples is 170. Figure 3.13 depicts this plot.



**Figure 3.13:** Froude nr vs Cavity closure angle detected

The detected cavity closure angle exhibits a decreasing trend with increasing Froude number. This behaviour is consistent with findings reported in the literature.

However, the uncertainty in the detected angles is not uniform across Froude numbers. The spread of the measurements varies considerably, and the distributions do not consistently follow a Gaussian profile. This suggests that the variability may not be solely attributable to random measurement noise, indicating that additional sources of systematic or regime-dependent uncertainty may be present. Further investigation is therefore required to identify the underlying cause of this behaviour.

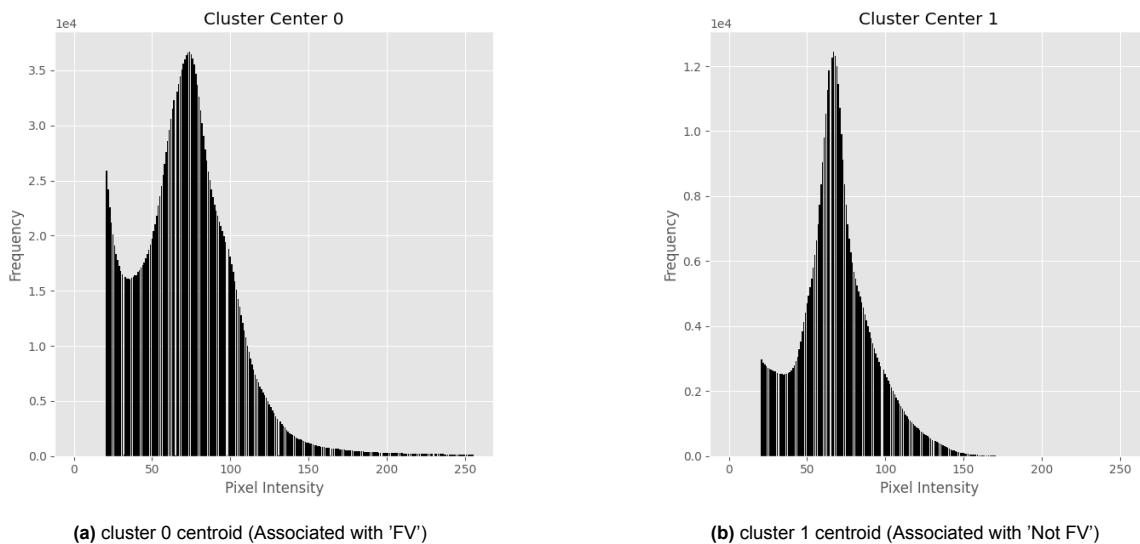
It is hypothesized that inaccuracies in the manual frame classification contribute to the increased variability observed. The magnitude of the variance is especially pronounced at  $F_n = 2.25$ . In particular, misclassification of transitional or partially ventilated states may lead to the inclusion of frames that do not exhibit a well-defined cavity geometry, thereby degrading the consistency of the closure angle estimation.

# Conclusion

In this thesis project, a methodology for the quasi-automated analysis and classification of flow features from visual data of a hydrofoil strut was developed. The applied methodology was split into three modules, namely Classification, Free Surface Tracking, and Cavity Closure Angle Reconstruction. The methods were developed on a specific dataset of towing tank images obtained at Delft University of Technology by Manuel Ferreira and Carlos Rodrigues [3].

Two iterations of a KMeans clustering algorithm were applied to attempt the classification of flow features, namely a  $k = 4$  clustering where a stratified random sample of flattened images were fed as the feature space. Stratified here indicating that the images were first containerized into their flow regimes based on human labels. It was decided to not accept these human labels as ground truth. It was found that although 3 out of 4 cluster centers have the appearance of a distinct flow regime, no reliable classification can be made from this clustering.


It was then decided to exploit the image artifact of a ventilated flow due to the air to water interface, namely a lens glare. A  $k = 2$  KMeans clustering algorithm, that takes the histogram of the grayscale image as its feature space, was trained on a specifically curated set of training data; a random sample with 50 % images of Fully Ventilated flows, and 50 % images of not fully ventilated flows. It was found that the cluster centers of this image are indeed different by their high frequency tails. This is a direct result of the lens glare that introduces many high intensity pixels in the center of the image.



**Figure 4.1:** Cluster centers from KMeans histogram clustering  $k = 2$

The performance of this  $k = 2$  histogram clustering was evaluated on a 200 image validation set with 20 fully ventilated images. The model has an accuracy of 76 % and correctly identified all True Positives .

The validation set was manually labeled 'FV' and 'Not FV'.



**Figure 4.2:** Confusion Matrix obtained from the validation data using the  $k = 2$  histogram clustering.

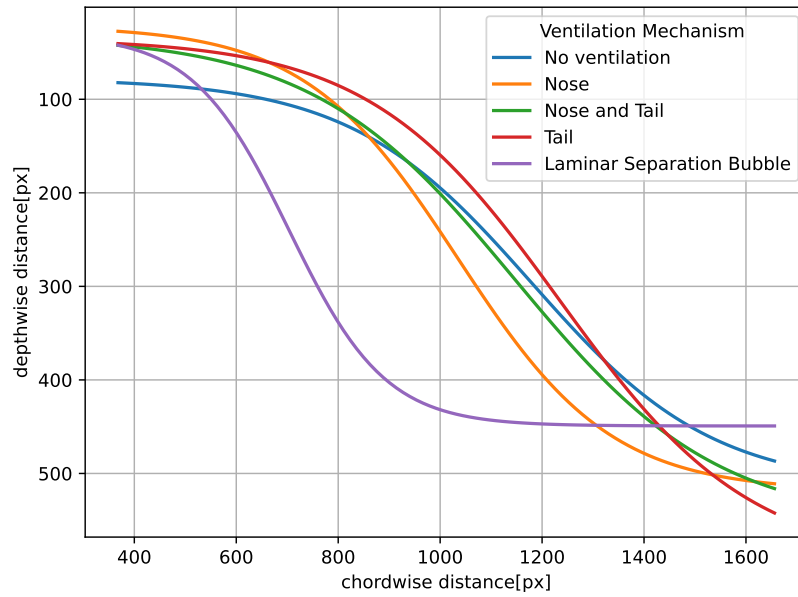
In addition to the clustering, two image processing pipelines were developed to extract spatio-temporal features of the flow from the video dataset. Namely the contour of the free surface of the water before the onset of ventilation, and the cavity closure angle after a fully ventilated cavity was developed.

An extensive stack of image pre-processing steps was employed to reconstruct the free surface or cavity closure angle of the flow from the image. Both methods are a modified version of the Hough Gradient Method for line segment detection, but with different parameters and different filters applied.

In the case of the free surface; background removal, Gaussian blurring, Canny edge detection, and the probabilistic Hough transform are applied in succession. This yields a set of line segments that can be combined into vertical lines, horizontal lines, oblique lines, and noise by the algorithm presented in Algorithm 1. The resulting vertical lines have their endpoints on the free surface such that the contour of the water can be fitted with an asymmetric sigmoid function. Applying this technique to every image in the dataset that is labeled as fully wetted results in 4021 free surface contours. Figure 3.9 displays the distribution of the error and Root Mean Square Error of these 4021 contours.

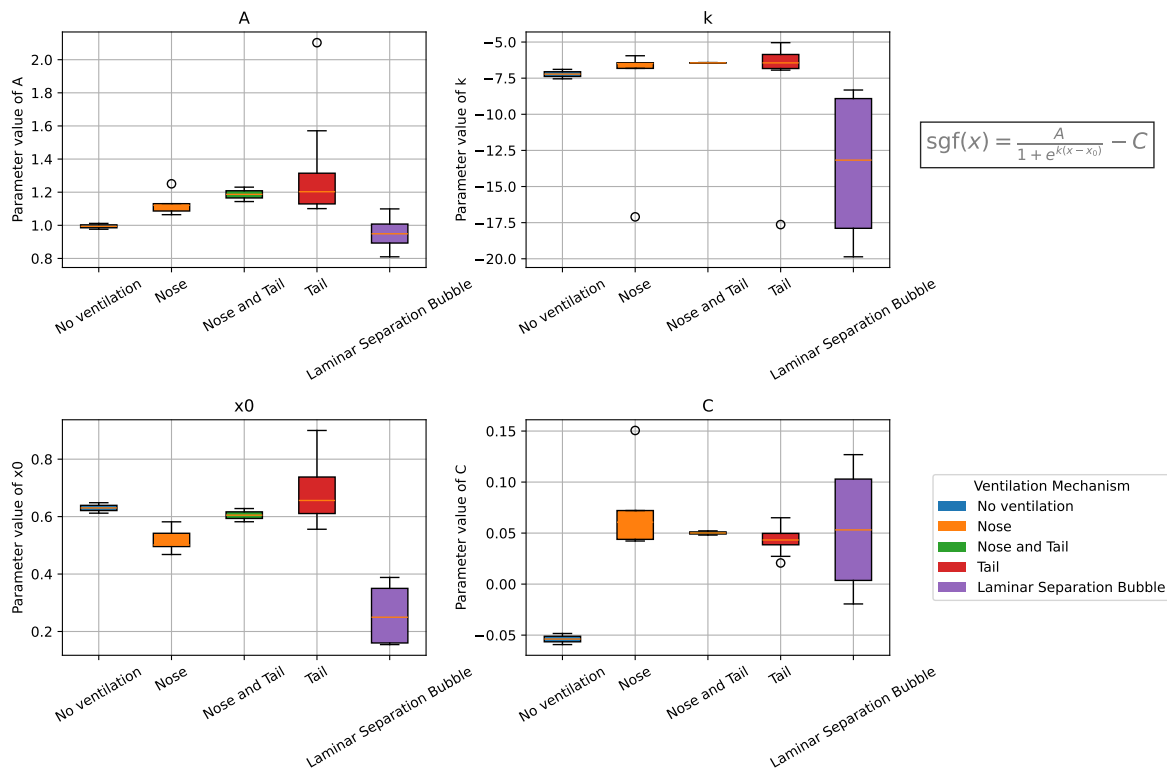
It was found that in 80 % of cases, the free surface detection results in a usable sigmoid function. In addition, this detection rate was found to be independent of the aspect ratio and dependent on the Froude number as illustrated in Section 3.2.

Averaging the usable sigmoid parameters on a per-video basis and clustering the result based on the observed ventilation mechanisms results in an average contour leading up to the different ventilation cases.



**Figure 4.3:** Per Ventilation mechanism mean asymmetric sigmoid curves for every frame before ventilation

It can be observed that the contour associated with the 'Laminar Separation Bubble' case looks different than the other ventilation mechanisms. A closer examination of the distribution of the individual parameters of these sigmoid function provides an insight in what is happening.




**Figure 4.4:** Per Ventilation mechanism mean of the parameters in the asymmetric sigmoid curve for every frame before ventilation

It can be seen that, compared to the other mechanisms,  $x_0$  is significantly smaller for the Laminar

Separation Bubble group. It was also observed that  $k$  is significantly more negative. This means an initial advancement has been made towards predicting the inception mechanism of a ventilated cavity from the free surface.

The algorithm for detecting the cavity closure angle is similar to the Hough Gradient Method for the Free Surface reconstruction. There is one major difference which leads to three additional pre-processing steps. Namely, the contrasts of the cavity closure angle are less sharp than those of the lines on the hydrofoil surface. As a consequence, No blurring was applied, the top half of the hydrofoil was masked such that the glare does not interfere with the detection, and the image was histogram equalized in order to amplify existing contrasts. After these steps, the Hough Gradient Method finds line segments in the vicinity of the cavity closure angle. A bandpass filter ranging from 5 to 45 degrees is applied to mitigate the noise and distinguish lines related to the cavity from lines related to the hydrofoil. The weighted average angle of the filtered line segments is reported as the cavity closure angle. Given the physical tendency of the cavity closure angle to get smaller with increasing Froude number as discussed in Section 3.3 and Section 1.2.1, the method is verified by plotting the distribution of the detected angle for each Froude number.



**Figure 4.5:** Froude nr vs Cavity closure angle detected

It was concluded that the method captures the physical trend of the cavity closure angle, but that a further investigation is needed to identify whether the distribution around the mean is, in some part, due to incorrect labeling of the image data, or due to the failing of the method.

## 4.1. Discussion

The methodologies developed in this work demonstrate that automated classification and geometric feature extraction from ventilated flow imagery are feasible within a controlled experimental setting. However, several limitations must be acknowledged.

First, the applicability of the developed framework is restricted by the experimental configuration. The classification and free surface detection methods were formulated and validated exclusively for the underwater, straight-on camera perspective, and only for the NACA 0010-34 geometry. Consequently, their robustness under different viewing angles, lighting conditions, or optical distortions remains unverified. In addition, the cavity closure angle extraction was only evaluated for the aspect ratio  $AR = 1$  configuration. Extension to other geometries, particularly  $AR = 1.5$ , would be very beneficial. For this specific research, that would require either a detection method that can detect the cavity

closure angle when the foil tip is outside the field of view of the camera, or a rerun of those experiments with a modified camera setup to include the tip of the deeper foil.

A further point of concern is the relatively large variance observed in the fitted parameters  $k$  and  $C$  of the sigmoid model used for free surface reconstruction. This variability suggests a potential inconsistency in phase boundary identification for this specific ventilation phenomenon. A modification to the algorithm to make the detected free surface points more robust would be a great addition to the current method, and could potentially foster a more conclusive distinction between the different ventilation mechanism.

More fundamentally, this work raises the question of how effectively a quasi-automated framework captures human interpretation of ventilated flow imagery. The free surface and cavity closure detection algorithms may be viewed as computational analogues of the manual boundary tracing performed by researchers such as Harwood and Young, who outline phase boundaries manually from their towing tank images. While the present approach introduces objectivity and reproducibility, it remains inherently dependent on thresholding, geometric assumptions, and fitting procedures that approximate what is, in practice, a visually interpreted boundary.

Finally, the attempt to leverage lens glare as a discriminative feature in the clustering procedure highlights an important methodological consideration. Although the glare artifact exhibited strong correlation with the Fully Ventilated flow regime, it does not originate from the intrinsic flow physics but rather from optical interaction at the phase boundary. As such, it cannot be considered a robust or generalizable feature, particularly for experiments conducted at greater submergence depths or under alternative lighting conditions. This underscores the importance of grounding classification strategies in physically meaningful features rather than experimental artifacts.

#### 4.1.1. Recommendations for Further Research

The results presented in this work demonstrate the feasibility of automated flow regime classification and cavity geometry extraction from high-speed imaging data. However, several limitations have been identified that provide clear directions for further research.

First, the unsupervised classification framework would benefit from more advanced machine learning techniques. While KMeans was able to distinguish 'FV' from 'Not FV', it cannot capture the complex interactions and nuanced differences between different flow regimes. Future work should investigate more powerful models. In particular, methods capable of handling non-linear separability and regime overlap are expected to improve robustness.

Second, the current free surface reconstruction approach captures the dominant interface geometry but does not resolve higher-order oscillations or local instabilities. Accurately detecting small-scale free surface fluctuations would significantly enhance the analysis of Rayleigh–Taylor instabilities and the Tail ventilation phenomena. Refinement of the detection pipeline is therefore recommended.

Third, the cavity closure angle extraction method requires systematic validation. A dedicated validation dataset, ideally including manually annotated ground-truth measurements, should be constructed. In addition, quantitative goodness-of-fit metrics must be defined to objectively assess the reliability and repeatability of the Hough Gradient Method. Such metrics would allow rigorous comparison between experimental conditions and processing configurations.

Beyond the cavity closure angle, additional feature extraction techniques should be developed to characterize ventilation inception mechanisms. Automated identification of air inception pockets, as well as the depth and contour of partial cavities, would provide a more complete geometric description of the ventilation process.

From an experimental perspective, improvements in image acquisition could substantially enhance data quality. The use of polarized lenses during towing tank experiments is recommended to reduce

lens glares, thereby improving feature detectability and reducing post-processing uncertainty.

Finally, I can see this field evolving towards volumetric and predictive approaches. Tomographic reconstruction of the ventilated cavity would enable three-dimensional characterization of the flow structure, overcoming the limitations of single-view imaging. A 3D reconstruction of the cavity would provide an excellent validation case for a numerical simulation of ventilated flows. Coupling such reconstructions with data-driven predictive algorithms may ultimately lead to models capable of anticipating ventilation onset based on evolving flow features. Such predictive models may one day inform the design and operation of high-performance sailing craft, improving ventilation control on primary control surfaces.

Together, these directions outline a pathway from the present two-dimensional image-based analysis toward a more comprehensive, robust, and predictive framework for hydrofoil ventilation research.

# References

- [1] A. J. Acosta. *Note on Partial Cavitation of Flat Plate Hydrofoils*. Technical Report E-19.9. Pasadena, CA: California Institute of Technology, 1955.
- [2] A. J. Acosta. "Hydrofoils and hydrofoil craft". In: (1973). DOI: 10.1146/annurev.fl.05.010173.001113. URL: <https://doi.org/10.1146/annurev.fl.05.010173.001113%7D>.
- [3] M. Aguiar Ferreira et al. *Ventilation of surface-piercing hydrofoils*. TU Delft - 4TU.ResearchData, Jan. 25, 2026. DOI: 10.4121/065ad4d4-224e-4984-a59d-6850db9dd1a8. URL: <https://doi.org/10.4121/065ad4d4-224e-4984-a59d-6850db9dd1a8..>
- [4] Manuel Aguiar Ferreira et al. "On the ventilation of surface-piercing hydrofoils under steady-state conditions". English. In: *Journal of Fluid Mechanics* 1028 (2026). ISSN: 0022-1120. DOI: 10.1017/jfm.2026.11126.
- [5] David Arthur and Sergei Vassilvitskii. "k-means++: The advantages of careful seeding". In: *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. SODA '07. New Orleans, Louisiana: Society for Industrial and Applied Mathematics, 2007, pp. 1027–1035. URL: <https://dl.acm.org/doi/10.5555/1283383.1283494>.
- [6] International Monohull Open Class Association. *IMOCA Class Rules 2025 V3.0*. Accessed: Feb. 4, 2025. Apr. 2023. URL: <https://www.imoca.org/mediacenter/uploads/2025-imoca-class-rules-v3-0.pdf?version=414e0203>.
- [7] G. Bradski. "The OpenCV Library". In: *Dr. Dobb's Journal of Software Tools* (2000).
- [8] C. E. Brennen. *Cavitation and Bubble Dynamics*. Oxford, UK: Oxford University Press, 1995. ISBN: 0-19-509409-3.
- [9] John P. Breslin and Richard Skalak. "Exploratory study of ventilated flows about yawed surface-piercing struts". In: 2-23-59W (Apr. 1959).
- [10] ARKEA Ultim Challenge Brest. *About the boats*. Accessed: Feb. 4, 2025. 2024. URL: <https://www.arkeaultimchallengebrest.com/en/page/the-boats-en>.
- [11] DutchSail. *JAJO Team DutchSail Women wint Europees Jeugd Kampioenschap 69F*. Accessed: 2025-03-06. 2024. URL: <https://www.dutchsail.com/nieuws/2024/7/22/jajo-team-dutchsail-women-wint-europees-jeugd-kampioenschap-69f>.
- [12] F. Feletti, E. Brymer, M. Bonato, et al. "Injuries and illnesses related to dinghy-sailing on hydrofoiling boats". In: *BMC Sports Science, Medicine and Rehabilitation* 13 (2021), p. 118. DOI: 10.1186/s13102-021-00343-8. URL: <https://doi.org/10.1186/s13102-021-00343-8>.
- [13] G. Franzosi et al. "Experimental study of bubble cavitation on a NACA 0015 hydrofoil by a computer vision approach". In: *Ocean Engineering* 316 (Jan. 2025). ISSN: 00298018. DOI: 10.1016/j.oceaneng.2024.119901.
- [14] Franzosi et al. "Analysis of Blade Root Cavitation Erosion with an Image Processing Approach". In: (2023). URL: <http://onepetro.org/ISOPEIOPEC/proceedings-pdf/ISOPE23/A11-ISOPE23/ISOPE-I-23-588/3164627/isope-i-23-588.pdf/1>.
- [15] Gibbs and Inc. Cox. *Hydrofoil Handbook, Volume 1: Design of Hydrofoil Craft*. Hydrofoil Research Project for Office of Naval Research, Navy Department, Washington, D.C., Contract No NONR-507(00). New York, NY: Bath Iron Works Corp., 1954.
- [16] C. M. Harwood. "The Hydrodynamic and Hydroelastic Responses of Rigid and Flexible Surface-Piercing Hydrofoils in Multi-Phase Flows". PhD thesis. University of Michigan, 2016.
- [17] C. M. Harwood. "The Hydrodynamic and Hydroelastic Responses of Rigid and Flexible Surface-Piercing Hydrofoils in Multi-Phase Flows". PhD dissertation. Ann Arbor, MI: The University of Michigan, 2016.

- [18] Casey M Harwood et al. “Experimental and Numerical Investigation of Ventilation Inception and Washout Mechanisms of a Surface-Piercing Hydrofoil”. In: 2014. URL: <https://www.researchgate.net/publication/280568075>.
- [19] Casey M. Harwood, Yin L. Young, and Steven L. Ceccio. “Ventilated cavities on a surface-piercing hydrofoil at moderate Froude numbers: Cavity formation, elimination and stability”. In: *Journal of Fluid Mechanics* 800 (Aug. 2016), pp. 5–56. ISSN: 14697645. DOI: 10.1017/jfm.2016.373.
- [20] Casey M. Harwood et al. “The hydroelastic response of a surface-piercing hydrofoil in multi-phase flows. Part 1. Passive hydroelasticity”. In: *Journal of Fluid Mechanics* 881 (Dec. 2019), pp. 313–364. ISSN: 14697645. DOI: 10.1017/jfm.2019.691.
- [21] Casey M. Harwood et al. “The hydroelastic response of a surface-piercing hydrofoil in multiphase flows. Part 2. Modal parameters and generalized fluid forces”. In: *Journal of Fluid Mechanics* 884 (2019). ISSN: 14697645. DOI: 10.1017/jfm.2019.871.
- [22] R. Huang et al. “Investigations into the ventilated cavities around a surface-piercing hydrofoil at high Froude numbers”. In: *Physics of Fluids* 34 (4 Apr. 2022). ISSN: 10897666. DOI: 10.1063/5.0085553.
- [23] G. Jacobi. “Pressure reconstruction on a T-foil from underwater multi-plane stereo PIV measurements in a towing tank”. In: *15th International Symposium on Particle Image Velocimetry – ISPIV 2023*. ISPIV. San Diego, California, USA, June 2023.
- [24] Y. Liu, Y. Wang, and K. An. “A feature extraction method for hydrofoil attached cavitation based on deep learning image semantic segmentation algorithm”. In: *Scientific Reports* 15.1 (Feb. 2025), p. 4415. DOI: 10.1038/s41598-025-88582-4.
- [25] Mackay Boats. *Bieker Moth*. Accessed: 2025-03-06. 2025. URL: <https://mackayboats.com/index.cfm/boats/bieker-moth/>.
- [26] J. Matas, C. Galambos, and J. Kittler. “Robust Detection of Lines Using the Progressive Probabilistic Hough Transform”. In: *Computer Vision and Image Understanding* 78.1 (2000), pp. 119–137. DOI: 10.1006/cviu.1999.0831.
- [27] K. I. Matveev, M. P. Wheeler, and T. Xing. “Numerical simulation of air ventilation and its suppression on inclined surface-piercing hydrofoils”. In: *Ocean Engineering* 175 (Mar. 2019), pp. 251–261. ISSN: 00298018. DOI: 10.1016/j.oceaneng.2019.02.040.
- [28] F. N. Nji et al. “Evaluation of Clustering Algorithms for Spatio-Temporal Multivariate Weather Data”. In: *Proceedings of the IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT)* (Dec. 6–9, 2022). Presented on September 20, 2022. Portland, OR, United States: IEEE/ACM, Dec. 6, 2022.
- [29] Nobuyuki Otsu. “A Threshold Selection Method from Gray-Level Histograms”. In: *IEEE Trans. Syst. Man Cybern.* 9 (1979), pp. 62–66. URL: <https://api.semanticscholar.org/CorpusID:15326934>.
- [30] C. N. Rodriguez. *Effects of a Quasi-Static Variation of the Angle of Attack in the Ventilation Inception of a Surface-Piercing Hydrofoil*. MSc. Thesis. Jan. 2025.
- [31] Peter J. Rousseeuw. “Silhouettes: a graphical aid to the interpretation and validation of cluster analysis”. In: *Journal of Computational and Applied Mathematics* 20 (1987), pp. 53–65. DOI: 10.1016/0377-0427(87)90125-7.
- [32] R. R. Schutt. “Unsteady Aerodynamics of Sailing Maneuvers and Kinetic Techniques”. PhD thesis. Ithaca, NY: Cornell University, Jan. 2017.
- [33] The Royal New Zealand Yacht Squadron. *THE PROTOCOL GOVERNING THE 37TH AMERICA’S CUP*. Accessed: Feb. 4, 2025. 2021. URL: <https://ac37noticeboard.officials.org/>.
- [34] R. Stigter et al. “An improved calibration methodology and uncertainty assessment in measurements of microbubble size and concentration”. In: *Experiments in Fluids* 66 (2025), p. 14. DOI: 10.1007/s00348-024-03929-3. URL: <https://doi.org/10.1007/s00348-024-03929-3>.

- [35] P. D. Swales et al. "The mechanism of ventilation inception on surface piercing foils". In: *Journal Mechanical Engineering Science* 16 (1 1974), pp. 18–24.
- [36] The HDF Group. *Hierarchical Data Format, version 5*. URL: <https://github.com/HDFGroup/hdf5>.
- [37] V. Vinayan and S. A. Kinnas. "A numerical nonlinear analysis of two-dimensional ventilating entry of surface-piercing hydrofoils with effects of gravity". In: *Journal of Fluid Mechanics* 658 (2010), pp. 383–408. ISSN: 14697645. DOI: 10.1017/S0022112010001783.
- [38] J. M. Wetzel. *Experimental Studies of Air Ventilation of Vertical, Semi-Submerged Bodies*. Technical Report Project Report no. 57. Minneapolis, MN: St. Anthony Falls Hydraulic Laboratory, University of Minnesota, July 1957.
- [39] Y. L. Young and S. Brizzolara. "Numerical and Physical Investigation of a Surface-Piercing Hydrofoil". In: *Third International Symposium on Marine Propulsors* (05 2013).
- [40] Y. L. Young et al. "Ventilation of Lifting Bodies: Review of the Physics and Discussion of Scaling Effects". In: *Applied Mechanics Reviews* 69 (1 Jan. 2017). ISSN: 00036900. DOI: 10.1115/1.4035360.
- [41] Y.L. Young et al. "Wave effects on the hydroelastic response of a surface-piercing hydrofoil. Part 2. Cavitating and ventilating flows". In: *Journal of Fluid Mechanics* 965 (June 2023). ISSN: 14697645. DOI: 10.1017/jfm.2023.254.