Machine learning and image analysis on photos of a solitary tree in a complex background- extraction and analysis of key properties for wind-tree interaction

Department of Wind Energy Master Report

Casper Chris Adriaan Bekkers

DTU Wind Energy-M-0349

July 2020

DTU Wind Energy
Department of Wind Energy

Author: Casper Chris Adriaan Bekkers

Title: Machine learning and image analysis on photos of a solitary tree in a complex background- extraction and analysis of key properties for wind-tree interaction

DTU Wind Energy-M-03499

July 2020

Project period:

November 2019 – July 2020

ECTS: 45

Education: Master of Science

Supervisors:

Ebba Dellwik

DTU Wind Energy

Anders Bjorholm Dahl

DTU Compute

Simon Watson

Axelle Viré

TU Delft

Remarks:

This report is submitted as partial fulfillment of the requirements for graduation in the above education at the Technical University of Denmark.

Technical University of Denmark
Department of Wind Energy
Frederiksborgvej 399
4000 Roskilde
Denmark

www.vindenergi.dtu.dk

Summary

Wind-tree interaction has been thoroughly studied both in wind-tunnels and in outdoor experiments. Some of this analysis has been performed by the use of image analysis. Though to the knowledge of the author, never over a long span of time or with a complex non-homogeneous background. This research presents a detailed description of a filtering and image processing technique developed for the specific case of analysing a single full-grown European Oak tree. The analysis spans over a time period of 6 months with a temporal resolution of 10 minutes. Another more specific analysis is performed over 10 days with a temporal resolution of 1 minute.

The developed method extracts enclosed surface area, tree surface area and horizontal center of area. By relating both surface areas to the wind speed, it is found that an increase in wind speed leads to a decrease in both surface areas as part of streamlining. Within this observation, a higher response for a full-leafed tree was detected. Furthermore, an increase in wind speed from a non-frequent wind direction leads to an inverse effect for lower wind speeds as an increase in surface area was observed. This indicates that the ability to efficiently streamline in one direction can come at the cost of the ability to streamline in another. With a comparison between the enclosed and tree surface area, a higher response of the tree surface area relative to the enclosed surface area is found. This indicates an increase in porosity with wind speed. A relationship between porosity and wind deficit is established and found to be non-linear, indicating the possibility of a high influence of skin-friction drag on general drag felt by the tree even for high wind speed regimes.

For the more specific 10 day investigation, a strain-gauge is used to acquire the drag force. This drag force is related to the wind speed to investigate the contribution of the changing surface area to the Vogel exponent. This contribution was found to be roughly half and thus indicates a high probability of another not yet quantified contribution to the drag force for wind-tree interaction.

Preface

This thesis is the final step in my graduation process for a combined double master degree of "Aerospace engineering" at the Delft University of Technology and for "Wind Energy" at the Technical University of Denmark.

The thesis is intended to be read by those interested in wind-tree interaction, but provides relevant background information in wind resource assessment, wind-tree interaction, image analysis and machine learning techniques used.

Lyngby campus, Kongens Lyngby, July 29, 2020

Casper Chris Adriaan Bekkers

Acknowledgements

Firstly, I would like to state that writing and developing this master thesis has been a truly inspiring and fun journey for me. I am fully aware however that I could not have performed this thesis by myself.

To this extent, I would first like to thank my supervisors Ebba Dellwik, Anders B. Dahl, Simon J. Watson and Axelle C. Viré for allowing me to ask for support if I ever had questions. Both Ebba and Anders, I would like to thank especially for their frequent involvement. Where Ebba helped me put the thesis in perspective, Anders had helped me surmount image analysis related obstacles and guided me to discover much needed background information in the field.

Furthermore, I would like to thank Nikolas Angelou for sitting in on most meetings with Ebba and providing me with helpful feedback related to the analysis and interpretation of the results.

Beyond help with the thesis, I would also like to thank my friends Charles Debusscher, Mario Garzón and Kamran Alimagham for working on their thesis besides me during the difficult time of the COVID-19 outbreak. They helped in keeping me inspired, happy and driven throughout the final steps for which I am truly grateful.

Finally, I would like to thank my family for their interest and their continued support of believing in me throughout my studies and life.

As this thesis is a final step in official education for me, I would like to cite a poem by Piet Hein, a Danish mathematician, scientist, inventor and poet, which I had found inspiring whenever encountering problems:

"The road to wisdom?—Well, it's plain
and simple to express:
Err
and err
and err again,
but less
and less
and less."
- Piet Hein

Contents

CHAPTER 1

Introduction

If specific information is available about the behaviour of wind, it can be used to determine where fragile structures should not be placed, where wind turbines can gather the most energy and how the weather will evolve. Already, for these reasons alone, researchers have been interested in getting a better grasp at the ever enigmatic wind. For wind resource assessment, sites are analysed and patterns are detected, which support thought processes that led to the discovery of relations. The relations in turn are used to construct models that allow, to some extent, the prediction of the wind.

As more and more turbines are placed, ranging up to 590 thousand turbines as of 2018 (GWEC, 2019), more sites need to be considered. The models in place right now (Troen and Lundtang Petersen, 1989) are shown to still contain erroneous results when it comes to vertical extrapolation (Boopathi et al., 2014), non-homogeneous terrain (Rathmann et al., 1996) and estimation near forests (Ebba Dellwik, Landberg, and Jensen, 2006). It is crucial for the development of these models, that more insight is gained on the interaction between the wind and the surroundings.

To this same goal, the interaction between forest lines and the wind is well investigated both numerically (RANS: e.g. Li, Lin, and Miller, 1990, LES: e.g. Cassiani, Katul, and Albertson, 2008) and analytically (e.g. Bache, 1986). More recently, single trees have also been subject to research, where mostly drag coefficients are linked to tree specimens as function of wind speed. Examples of this are the investigation of a red maple by Kane and Smiley, 2006, poplar trees by Koizumi et al., 2010 and a more general attempt has been made by Cao, Tamura, and Yoshida, 2012 who separately considered deciduous, coniferous and broad-leaf evergreen trees.

These investigations pertain in some way to the momentum transfer between the wind and the tree. This has been measured by either the bending moment experienced by the tree, or the velocity deficit measured near the tree. Momentum deficit however, also manifests itself to some degree in optically inspectable parameters of the tree such as the optical porosity, deformation or reshaping of the crown as simply hinted at by the Beaufort scale.

One specific example of a single tree investigation is that of a near wake measurement of a single full-grown European Oak tree compared to a RANS model by E. Dellwik et al., 2019. This research has been conducted as part of a research project launched by the "Danmarks Frie Forskningfond" which has been launched in August 2016 and will run until July 2020[1] and aims to investigate the interaction of singular trees with near-surface

[1] https://www.vindenergi.dtu.dk/english/research/research-projects/single-tree

winds. As part of the project, two meteorological masts with sonic anemometers have been set-up near the tree. More important for the research in this thesis though, is the inclusion of a camera that has been taking pictures of the tree, every minute for the last two years, collecting large amounts of data.

The aim of this research is two-fold. Below these two aims are indicated:

- **Examine to which extent key parameters for wind-tree interaction can be extracted from images taken of a tree in complex circumstances with a non-homogeneous background.** The extend to which this can be achieved will be described by whether concurrence is found between found and expected interaction. These expected interaction are the decrease of surface area with wind speed, swaying of the tree based on wind direction and wind speed, the increase in porosity with wind speed and a higher surface area response in periods with more leaves.
- **Explore the capability of these key parameters to better describe the interaction between the tree and the wind.** Interactions explored are those of the changing surface area based on wind direction and speed, the exact relationship between wind speed and porosity and the effect of surface area on the drag over wind speed curve.

These aims will be achieved by establishing background knowledge and the research conditions such as the experimental set-up. Furthermore, a methodology is created which filters out unusable images and systematically extracts image aspects from left-over images. The results of this filtering process and the linking between these image aspects and wind statistics are stated and discussed.

1.1 Background and state-of-the-art

Four main themes will be central in the discussion of this thesis. These themes need to be described sufficiently before anything substantial can be said about the set-up, the methodology and the results of this thesis. This section aims to provide the reader with relevant background knowledge on the central subjects and present their state-of-the-art as well as describe central techniques used in the development of the methodology.

The four subjects that are central in this thesis are wind resource assessment, wind-tree interaction, image analysis and machine learning. Wind resource assessment is required for gathering and understanding wind statistics that are to be linked with the image aspects. Wind-tree interaction is the field in which this thesis provides new insights and uses current knowledge as means of verifying observations. Image analysis is the field that is used to extract the image aspects that will eventually be linked to wind statistics. It is crucial in processing images into tangible information. Finally, machine learning will be used as a tool throughout the thesis, both for training classification models and to investigate, highlight and describe relations between statistics.

1.1.1 Wind resource assessment

The wind is interesting for many stakeholders. It can indicate weather, estimate loads on structures and provide insights on performance changes in air and water based transportation. From an energy production standpoint however, it is most interesting for having kinetic energy that can be extracted. Wind resource assessment is critical for the decision of location, estimation of structural loads and prediction of performance of wind turbines. For these reasons, an accurate assessment of the wind is highly sought after in the wind industry.

The features of the wind that are extracted are usually speed, direction and in some cases inclination angle. These features are extracted by transducers at a high sampling frequency. This leads to high amounts of data. To better store and process this data, they are transferred into statistics.

Statistics

A good way to reduce the size of data is to reduce the temporal resolution. For this reason, wind statistics are often recorded in 10-minute averages. When time-averaging wind direction, the vector-wise mean is used to take the jump from 360° to 0° into account. If only the mean of a signal is kept however, it is evident that a lot of information is lost. For instance, an arbitrary signal that starts around 0 and with high fluctuations moves up to 10, will be recorded as 5, where the same is true for a signal that stays steadily around 5 from beginning to end. To mediate this loss of information, besides the mean, the standard deviation and in some cases also the minimum and maximum of the signal are recorded.

To get an even better profile of the wind beyond a signal $U(t)$, a probability density function $f_X(U)$ is fit to a wind-speed histogram of the signal. The distribution described by Weibull, 1951, known as the Weibull distribution, is most commonly used for this purpose as it fits most every 10-minute averaged wind-speed histogram well. The distribution is shown in 1.1 where k is the shape parameter and λ is the scale parameter in meters per second. The fit of this function to an arbitrary time series of a wind signal is shown in Figure 1.1.

$$f(x;k,\lambda) = \begin{cases} \frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k} & x \geq 0 \\ 0 & x < 0 \end{cases} \tag{1.1}$$

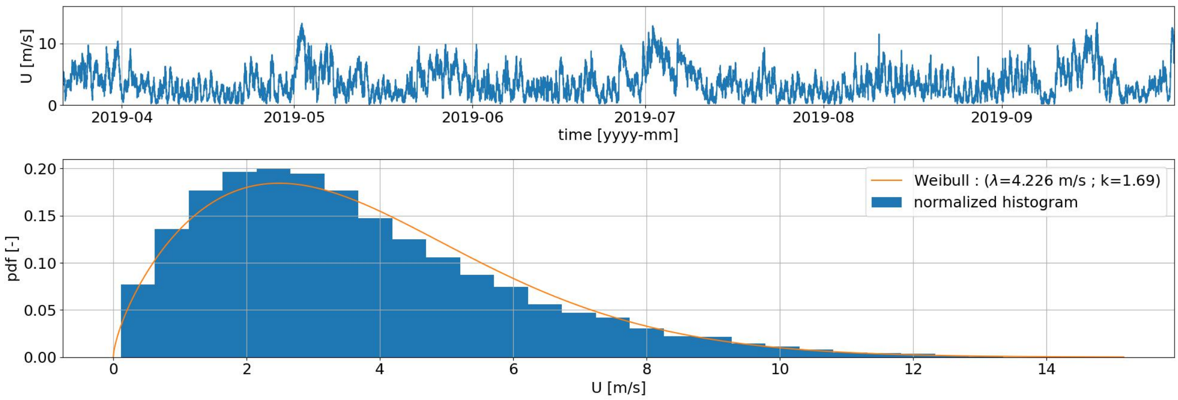

Figure 1.1: Example of calculating the probability density function of a time-signal of 10-minute averaged wind speeds.

With the mean $\bar{U}$ and the variance σ_U^2 of the wind speed, the turbulence intensity U_{TI}, which is highly indicative of the amount of turbulence in the air, can be calculated. The turbulence intensity is calculated by the relation shown in Equation (1.2).

$$U_{TI} = \frac{\sigma_U}{\overline{U}} \tag{1.2}$$

Sensors

In order to gather wind statistics, signals need to be collected with the help of transducers. When talking about wind speed and wind direction, it is most conventional to use either wind vanes combined with cup anemometers, sonic anemometers or LIDAR (Light Detection And Ranging). Of these three, sonic anemometers will be explained as they have been used for the purposes of this research.

Transducers are a way of measuring the influence of a physical phenomena that can be linked to a parameter of interest. In the case of a sonic anemometer, the observation that sound travels quicker when its medium is moving in the same direction is used.

An ultra-sonic pulse is send from one end of a transducer pair to the other and back. The difference in time it takes between the first and the second pulse is indicative of the speed along the axis of transmitter and receiver V_l. A schematic of a sonic anemometer setup and the relevant parameters is displayed in Figure 1.2 where V_∞ is the ambient wind speed and a is the speed of sound.

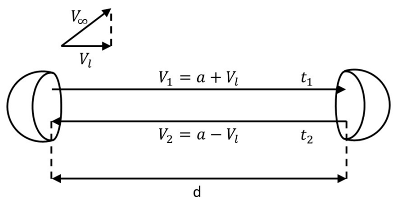

Figure 1.2: Schematic of a sonic anemometer depicting its relevant variables.

By using the relationship between speed and travelled time shown in Equation (1.3), the tangential wind speed can be calculated as seen in Equation (1.4).

$$t_1 = \frac{d}{V_1} = \frac{d}{a + V_l} \qquad t_2 = \frac{d}{V_2} = \frac{d}{a - V_l} \tag{1.3}$$

$$V_l = \frac{d}{2}\left(\frac{1}{t_1} - \frac{1}{t_2}\right) \tag{1.4}$$

This means that the magnitude and direction of the wind speed along the line of a transducer pair is determined by t_1, t_2 and d. For most sonic anemometers three of these pairs are positioned such that their lines of sight do not share a 2D plane. This means that by using trigonometric relationships the wind speed is known in any direction.

When the distance is roughly 20 cm, the ambient temperature is 20 degrees and there is no wind, this means that a pulse can be send more than 800 times every second. However, as the spatial resolution of the anemometer is as high as the distance between the transducers, having a temporal resolution of more than 800 Hz would be dis-representative. In general a sampling frequency of 20 Hz is used.

It should be noted that there is physical interaction between the wind and the sonic anemometer. Using only the analytical model would thus result in incorrect measurements. To correct this, a calibration process is performed to ensure that any RMS error is below 1%. Such a calibration process from laminar wind tunnel conditions to turbulent atmospheric conditions has been evaluated by Högström and Smedman, 2004. It was found that whereas the calibration process had decreased the average error in the tunnel to a range of 0.4-0.8%, in atmospheric conditions this error was wind-direction sensitive and was in the range of 5%. In a newer study (Peña, Ebba Dellwik, and Mann, 2019) a method of assessing the accuracy of the sonic anemometer was proposed by looking at the ratio between the vertical and horizontal power spectrum (described to be 4/3 by Kolmogorov, 1941) of the inertial subrange. It had been found that the "Metek GmbH" anemometer's calibration process had retained the ratio well, which is used in this research.

1.1.2 Wind-tree interaction

Wind-tree interaction highly considers fluid structure interaction (FSI) between the wind (fluid) and the tree (structure). However, the field of wind-tree interaction behaves differently from standard FSI applications. Man made structures are made for strength and stiffness as large deformations are usually undesirable. Contrarily, plants have no dire need to be stiff. This means that they are mainly made for strength and deformation occurs readily (Denny and Gaylord, 2002).

Energy (kinetic and thermal in nature) is transferred between the wind and the tree when the wind blows past the tree. This interaction happens on different levels. Mostly, a drag force will be exerted on the wind by the tree which will reduce its kinetic energy. A drag force F_D of a bluff body for high enough Reynolds numbers is described by Hoerner, 1965 as influenced by an empirically determined drag coefficient C_D, a frontal surface area S, the free stream velocity U_∞ and density of the fluid ρ. This relation is shown in Equation (1.5) where it is noted that if shape, drag coefficient and air density are constant, the drag force will change proportional to velocity squared.

$$F_D = \frac{1}{2}\rho U_\infty^2 S C_D \qquad \rightarrow \qquad F_D \propto U_\infty^2 \tag{1.5}$$

However, when looking at the drag force on a tree with an increasing wind, it does not increase proportional to the velocity squared as most bluff bodies do (for example Kane and Smiley, 2006, Koizumi et al., 2010 and Cao, Tamura, and Yoshida, 2012). A further study on the mechanics of a plant in fluid flow by Gosselin, 2019 explains this by applying a dimensional analysis to the problem and concludes three reasons. Firstly, due to the streamlining of the leaves and the tree, the drag coefficient reduces. Secondly, as the tree presents itself as a porous structure, it slows down the air that passes through it. Finally, also due to the streamlining, the frontal surface area of the tree changes. This means that C_D, U and S are all functions of the wind speed.

This concept has been described more empirically by Vogel, 1984. Vogel has done research on the speed specific drag F_D/U_∞^2 versus speed U_∞ for a loblolly pine and a holly. Here the speed specific drag saw a peak at around 5 and 8 m/s and then tapered off following a different exponent. This led to the inclusion of the so called Vogel exponent E when describing the drag force and wind speed relationship of a tree. For higher speeds and for a specific tree, it is thus expected that the drag follows Equation (1.6) with E being a negative number.

$$F_D \propto U_\infty^{2+E} \tag{1.6}$$

1.1.3 Image analysis

A digital image is an array which describes the value of pixels. Three types of images are discussed, ranging from simple to complex they are: binary, gray-scale and color. A binary image is also known as black and white. Each pixel will be either a one (white) or a zero (black), making it a simple method of storing masks and divisions of pixels into two classes. One step higher in complexity is a gray-scale image. A gray-scale images states for each pixel where its value is on a scale ranging from black (0) to white (255). This allows for a resolution of 2^8 colors. Finally, the most complex image, is a colored image with usually represented with RGB-channels (Red Blue Green). Here each pixel has three values ranging from 0 to 255 describing the presence of red, green and blue in the image respectively, leading to a posibility of 2^{24} colors. A representation of a color image is shown in Figure 1.3.

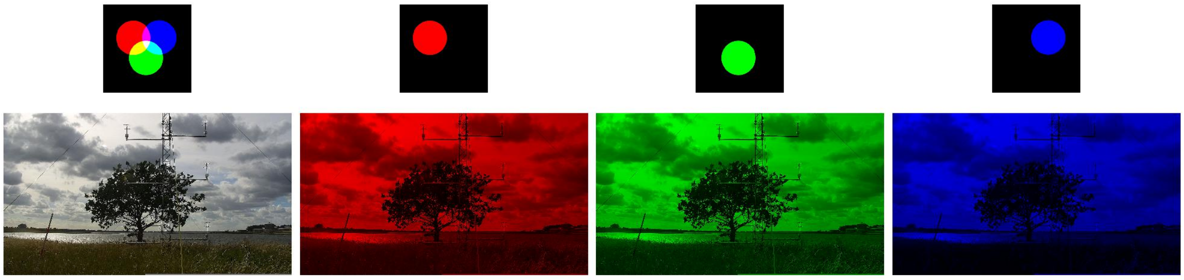

Figure 1.3: Decomposition of a colored image (left) into its red (R), green (G) and blue (B) channels.

Four image analysis concepts will be readily used during this research. These are thresholding, Gaussian filters, down-scaling and image gradients.

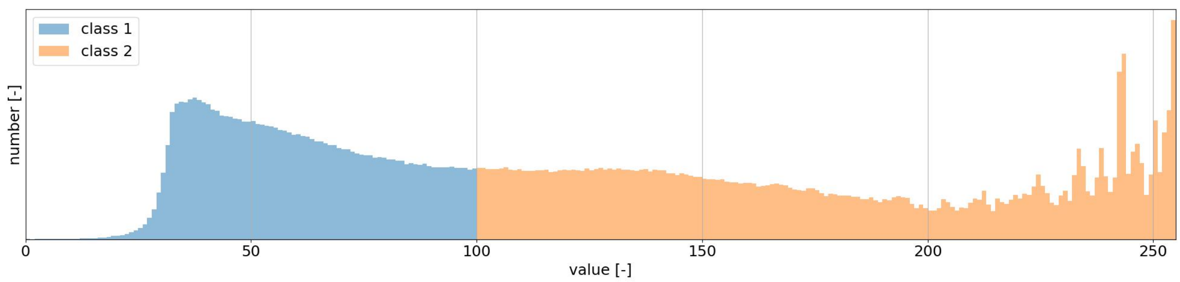

Figure 1.4: Visualisation of thresholding a histogram with a threshold set at a value of 100.

Thresholding

A method of binarizing a single channel image is thresholding. Thresholding in its most general sense means setting a specific value and classifying values below it to one class and those above it to another. An arbitrary example of this classification is shown in Figure 1.4.

For image analysis purposes a histogram can be made from an RGB-channel or HSL-channel (Hue, Saturation and Luminosity). The threshold can be set manually to select which pixels will belong to a specific class. For large amounts of images however, this process will quickly become tedious and an automatic method of selecting a threshold is desired. In a study by Mizoue and Inoue, 2001 three different automatic thresholding methods are discussed in the scope of separating trees from their background. By inspecting three different statistics of a histogram, a threshold can be set for an arbitrary image. The three different thresholding techniques that are discussed are those of maximizing between-class variance by N., 1979, minimizing classification error by Kittler and Illingworth, 1986 and maximizing entropy by Kapur, Sahoo, and Wong, 1985.

Calculating these three values, requires four statistics based on the histogram with bins $[1, 2, ..., L]$ where each i'th bin contains p_i fraction of the total pixels. The bins are split at a threshold k, leading to statistics shown in Equation (1.7). Where ω_j represents the fraction of pixels that belong to the class j, μ_0 represents the mean of class j, σ_0^2 represents the variance of class j and E_j represents the entropy of class j.

$$
\begin{aligned}
\omega_0 &= \sum_{i=1}^{k} p_i & \omega_1 &= \sum_{i=k+1}^{L} p_i \\
\mu_0 &= \sum_{i=1}^{k} i p_i / \omega_0 & \mu_1 &= \sum_{i=k+1}^{L} i p_i / \omega_0 \\
\sigma_0^2 &= \sum_{i=1}^{k} (i - \mu_0)^2 p_i / \omega_0 & \sigma_1^2 &= \sum_{i=k+1}^{L} (i - \mu_1)^2 p_i / \omega_1 \\
E_0 &= -\sum_{i=1}^{k} \frac{p_i}{\omega_0} \log \frac{p_i}{\omega_0} & E_1 &= -\sum_{i=k+1}^{L} \frac{p_i}{\omega_1} \log \frac{p_i}{\omega_1}
\end{aligned}
\tag{1.7}
$$

These statistics can be used as described by N., 1979, Kittler and Illingworth, 1986 and Kapur, Sahoo, and Wong, 1985 to obtain the between-class variance σ_B^2, the classification error J and the total entropy Ψ respectively as seen in Equation (1.8).

$$
\begin{aligned}
\sigma_B^2(k) &= \omega_0 \omega_1 (\mu_0 - \mu_1)^2 \\
J(k) &= \omega_0 \log \left(\frac{\sigma_0}{\omega_0} \right) + \omega_1 \log \left(\frac{\sigma_1}{\omega_1} \right) \\
\Psi(k) &= E_0 + E_1
\end{aligned}
\tag{1.8}
$$

Calculating the statistics for integer splits for a histogram of the blue channel of an image of a tree, leads to the graphs seen in Figure 1.5.

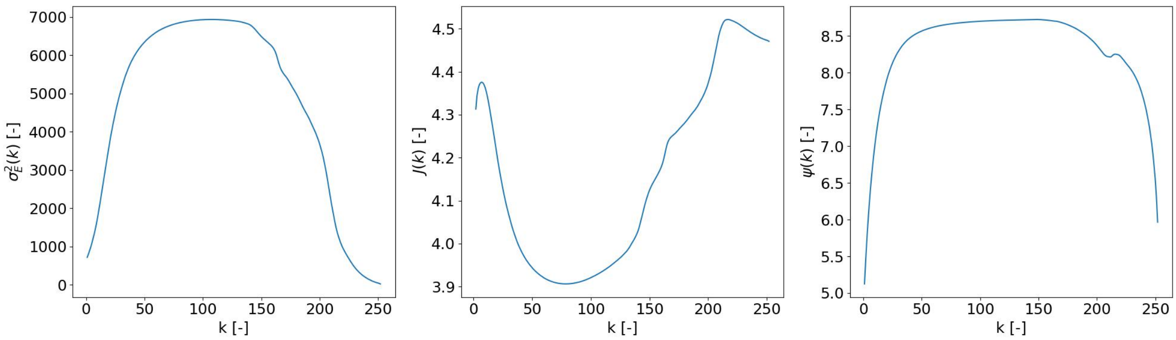

Figure 1.5: The between-class variance (σ_B^2), classification error J and total entropy Ψ for integer splits (k) between 1 and 254.

Applying the optimal split to the histogram of a blue-filtered gray-scale image of a tree for all three statistics leads to the three binarized images seen in Figure 1.6.

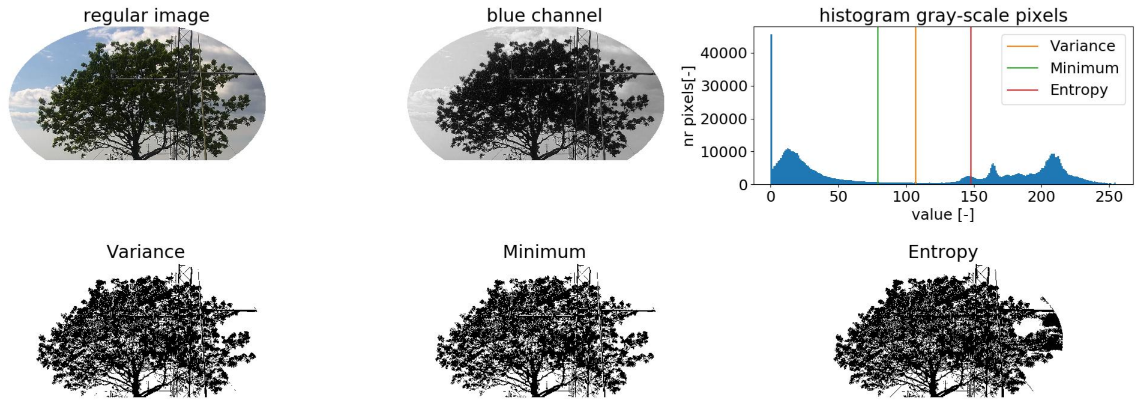

Figure 1.6: Result of thresholding the histogram (top right) of the blue channel (top middle) of an image (top left) with automatic thresholding based on maximum between-class variance (bottom left), minimum classification error (bottom middle) and maximizing entropy (bottom right).

Based on the application any of the three methods might be better suited. For trees however, it has been stated by Mizoue and Inoue, 2001 that variance thresholding corresponds most to interactive (manual) thresholding. Lee, Alfaro, and Van Sickle, 1983 has found that choosing the blue channel leads to the best ability to threshold images of trees with sky.

Gaussian filter

Convolving an image with a kernel is a way of modifying that image. This can be used for edge detection or to smooth an image. Smoothing in the simplest manner could be done by averaging a pixel with its neighbors. Gaussian filtering takes this process but averages in a weighted manner which follows a Gaussian distribution. Such a kernel would look like the one illustrated in Figure 1.7. It is important to consider how big to make a kernel for a specific Gaussian filter. In general it is a good rule to have a kernel size that is $6\sigma + 1$ pixels in each dimension.

0.0	0.0	0.001	0.002	0.001	0.0	0.0
0.0	0.003	0.013	0.022	0.013	0.003	0.0
0.001	0.013	0.059	0.097	0.059	0.013	0.001
0.002	0.022	0.097	0.159	0.097	0.022	0.002
0.001	0.013	0.059	0.097	0.059	0.013	0.001
0.0	0.003	0.013	0.022	0.013	0.003	0.0
0.0	0.0	0.001	0.002	0.001	0.0	0.0

Figure 1.7: Gaussian kernel with a standard deviation (σ) of one pixel.

Down-scaling

The method used for down-scaling a binary image in this report consists of three steps.

The first step is to divide the grid into $r \times r$ cells. Each of these cells will eventually become a pixel for a down-scaled image. The second step is to set a threshold for the minimum amount of pixels required to be present in a cell for it to become an active pixel. Finally, any cell with pixels above this threshold will be counted as an active pixel in the new down-scaled image. An example of this with a lenient threshold is seen in Figure 1.8.

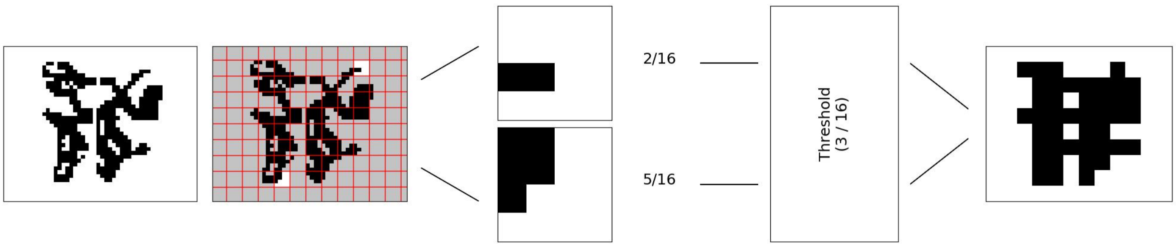

Figure 1.8: Visualisation of down-scaling an arbitrary image with a value of $r = 4$ and a threshold of 3/16. Showing from left to right, the binary image, the subdivision into cells, two examples of these cells, the threshold and the down-scaled image.

Image gradient

For purposes of extracting the gradient from an image, the image can be seen as a 2-D surface. Which means that every point (pixel) will have a gradient in two directions: horizontal (x) and vertical (y). These gradients are calculated through finite differences. Many finite differences exist, but for computational efficiency, the easiest to use is forward or backward finite difference. Forward finite difference in the horizontal and vertical direction is seen in Equation (1.9) and (1.10) respectively. With this, the gradient in the image at coordinate i (in the x-direction) and j (in the y-direction) can be calculated in the form of a change in value I with a pixel distance d_p. Note that the first coordinate represents the y-position (row of the image) and the second coordinate represents the x position (column of the image).

$$\frac{\delta I(j,i)}{\delta x} = \frac{I(j,i+1) - I(j,i)}{d_p} \tag{1.9}$$

$$\frac{\delta I(j,i)}{\delta y} = \frac{I(j+1,i) - I(j,i)}{d_p} \tag{1.10}$$

1.1.4 Machine learning

Artificial intelligence was first discussed by Turing, 1950 where the question was posed whether a machine could posses human intelligence. A part of this is to program a machine to have intelligence through if-then statements and hard-build rules. However, another route can be taken in providing a machine with intelligence. Giving it the ability to learn from experience. This second method is a sub-set of AI: machine learning, first applied by Samuel, 1959. It is mostly based on statistical models that improve with processing data (experience). These statistical models can be interpretable and can work exploratory towards data. However, understanding of how the machine connects data and finds relations can be lost, e.g. it becomes a black box. This usually happens as soon as neural networks are used with multiple layers. This part of machine learning is called deep learning. Thus, deep learning is a subset of machine learning which in turn is a subset of artificial intelligence.

In broad terms, the functionality of machine learning is stated well by T. M. (M. Mitchell, 1997 as: "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.". To get a good view on what machine learning is, these three variables (experience, task and performance) need to be discussed.

For this thesis, the practices used will remain in the interpretable part of machine learning to explore data and find relations (no use of multi-layered neural networks). Before more can be stated about the state-of-the-art of machine learning, some nomenclature needs to be established which will also explain the form of the data or *experience*. When this is established, the purposes of the machine learning techniques used in this report or the *tasks* are discussed. These purposes (tasks) are initial inspection of the distribution of data, finding relations between statistics, detect outliers in the data and predicting statistics through classification. Finally, methods of validating and evaluating models that are fit to the data are discussed. In other words, the *performance* measure of the model.

Nomenclature

Machine learning revolves around processing data or *experience*. Data can come in many forms, thus, to create a consistency in describing it, two terms are used: observation and attribute. An observation describes an instance, while the attributes describe the characteristics of an observation. The documented attributes will all lie along an axis of interest. For example, 10 observations can be made of different images, where each has different attributes for surface area and center of area. Within any dataset, the number of instances will be indicated with N and the number of attributes will be indicated with M.

In machine learning, a distinction can be made on the type of learning. This can be either supervised or unsupervised. Supervised learning is that of a feedback loop that is used to fit a model. The model will make a prediction, and the difference between its output and the actual output will improve the model. Unsupervised learning, is that of the model not knowing the actual output, and still trying to improve. This may seem impossible, but methods such as clustering (first described at a rudimentary level by Driver and Kroeber, 1932) and association rule mining by Agrawal and Srikant, 2013, are tuned to find patterns in the data.

A final note is that of regularization. Many fitting techniques are prone to over-fitting. This leads to a lowered cost function on the training set, but a higher one on a test-set (unseen data). To prevent this, model complexity parameters are weighted to some extent and added to a performance measure (or cost function) that is to be minimized. This penalizes the use complexity parameters and suppresses over-fitting. When there is a high discrepancy between training and test performance, regularization is looked into.

Inspecting data

Understanding the basic shape of a data-set is not only useful for verifying your data, gaining an intuition of its behaviour and making some predictions as to which relations are present, it is also a valuable tool for initial verification of results in further processes. Inspecting data is a first step towards understanding the relations that occur within it.

While it is easy to make sense of a single instance with attributes. As the number of instances increases, this becomes harder. Two important insights are made by inspecting data. First, the distribution of an attribute can give insights on the relative position of an observation in a data-set. It will also point out odd observations (usually linked to outliers) when the relative position of an attribute is highly unlikely given a distribution of other data points. Second, visually inspecting the position and clustering of observations relative to each other gives insight on correlations, and tendencies between attributes.

Relate attributes

Linear relationships are best investigated by looking at their correlation or co-variance. When relationships become nonlinear however, a good solution is to assume a shape of their relationship based on an initial investigation and fit a curve to it to minimize the summation of the squared error of each point to the fit. This is done through the algorithm described by Levenberg, 1944.

Another way to relate attributes is to find a moving average of the attributes plotted against each other. This can be a good descriptor to compare any fit with and give a general tendency of the data. For these moving averages, two types are considered. The simple centered moving average, which consists of setting a points value equal to the mean of N points around it, where N is a value chosen based on the level of smoothing required. A high value of N smooths a signal more, at the cost of loosing a sharp step response. A low value of N smooths a signal less with the benefit of a high step response. When interested in the frequencies in a signal however, the simple moving average does not work well as described by Smith, 1997. In these cases, although coming at a high cost of computational effort, Gaussian smoothing is better suited as a low-pass filter.

Detect outliers

Besides creating models, machine learning is used for outlier detection. This will be relevant for both the cases where a picture has been taken that does not correspond to expectations (a bird flies in front of the tree, rain drops on the lens, etc.) and those where the meteorological mast has erroneous measurements. This outlier detection can be done at early stages by inspecting the box plots and finding points that clearly do not follow the distribution, or by density estimations as described by Fix and Hodges, 1951. Two examples of these are kernel density estimator proposed by Rosenblatt, 1956 and average relative density described by Tan et al., 2019. This principle gives a density score to each observation and if below a chosen threshold, a sample will be regarded as an outlier and thus not considered for further analysis.

For this analysis however, filters are used instead of outlier detection. After this filtering process, outlier detection is seen to be redundant.

Classification

Logistic regression is a method of applying regression for classification first discussed in Garnier and Quételet, 1838. It does so, by fitting a model that takes attribute values and outputs the probability (between 0 and 1) that the observation belongs to a positive class. This probability function θ takes to form as seen in Equation (1.12), which is visualized in Figure 1.9. In this function, a_i^j is the j'th attribute of the i'th observation, w_j is the j'th weight, b is the bias and $\sigma(x)$ is the Sigmoid function shown in Equation (1.11).

$$\sigma(x) = \frac{1}{1+e^{-x}} \tag{1.11}$$

$$\theta(\boldsymbol{a}_i, \boldsymbol{w}) = \sigma\left(b + \sum_{j=0}^{M} a_i^j \cdot w_j\right) \tag{1.12}$$

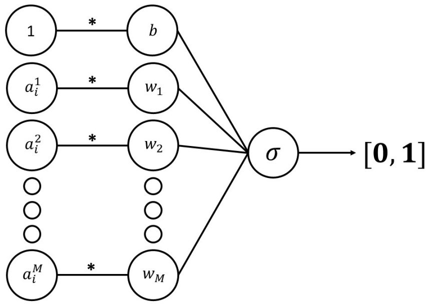

Figure 1.9: Visualisation of logistics regression. Attributes and a bias are multiplied by a weight and combined into the Sigmoid function to produce a value between zero and one.

Combining the probability in Equation (1.12) with the Bernoulli distribution, leads the the probability density function shown in Equation (1.13).

$$p(y_i|\boldsymbol{a}_i, \boldsymbol{w}) = \theta^{y_i} * (1 - \theta)^{1-y_i} \tag{1.13}$$

The aim of a logistic regression is to fit weights such that a cost function is minimized. A cost function for a logistic regression is shown in Equation (1.14). In this equation the final term is the regularization term with λ as the regularization parameter.

$$E(\boldsymbol{w}) = -\frac{1}{N} log \left(\Pi_{i=0}^{N} p(y_i|\boldsymbol{a}_i, \boldsymbol{w}) \right) + \lambda \left\| \boldsymbol{w} \right\|^2 \tag{1.14}$$

This equation does not have an exact solution and some sort of optimization is required to minimize the cost function. For this purpose, limited-memory BFGS optimization described by Liu and Nocedal, 1989 will be applied.

Feature selection

When choosing which attributes will be used for a classification or regression model, it might not be the best idea to include as many attributes as possible. More attributes, allow the model to over-fit the data more, which will increase the generalization error (a generalization error is the error that is expected on an unseen dataset). However, when multiple attributes are available, investigating model performance on all permutations of attributes to see which one will perform best will be computationally expensive or even infeasible. For this reason, feature selection can be applied. Two versions of feature selection are highlighted: forward and backward selection.

Consider a model with M possible features. Forward selection starts with a model which uses zero features, this will act as a baseline, and will have an accuracy as big as the fraction of observations in the biggest class in case of classification or variance in case of regression. Next, M models are made which include all previous features (in the initial case none) and one of all features that have not yet been used. These models are evaluated individually, and the model with the best performance is used as the next baseline. After this, the process is repeated with $M-1$ new models, where each new model includes the features of the best performing previous model and one of the remaining attributes. This process is repeated until adding a new feature will no longer improve the model, in which case the selection has converged.

Backwards selection starts with a model which uses all M features as a baseline. Then M new models are made, where one feature is left out. The model of these that performs best, is then selected as the next baseline model containing $M-1$ features. Next, $M-1$ new models are made, where one extra feature of the previously best performing model is left out. This process is once more repeated until removing a feature, no longer increases the performance of the model.

Performance measure

In regression and classification, some sort of error measure is minimized. For regression this will often be the mean squared error. In classification, either accuracy, or precision given a minimum recall will be optimized. To understand precision, recall and accuracy, a confusion matrix is shown in Figure 1.10 to visually introduce true positives (TP), false positives (FP), true negatives (TN) and false negatives (FN).

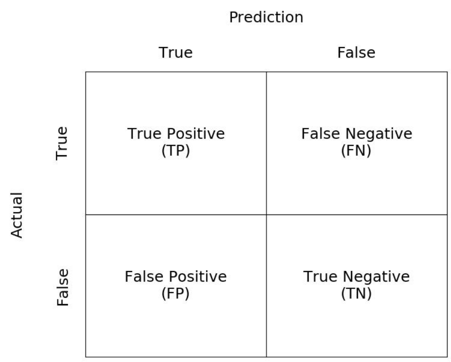

Figure 1.10: Visual representation of a confusion matrix.

Recall, precision and accuracy are some combination of these four classes. Recall is the fraction of positives that have also been predicted as positive. In other words, the number of true positives over the total number of positives:

$$\text{Recall} = \frac{TP}{TP + FN} \tag{1.15}$$

Precision is the fraction of total predicted positives that are correct. In other words, the number of true positives over the number of total number of positive predictions:

$$\text{Precision} = \frac{TP}{TP + FP} \tag{1.16}$$

Accuracy is the fraction of how many predictions were correct. In other words, the number of true positives and the number of true negatives over the size of the dataset:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{1.17}$$

A classification model is trained to minimize either the accuracy or precision given a minimum recall. Training a model on all available data however, does not state the performance on unseen data. A generalization error is to be estimated. To be able to measure this, unseen data can be artificially provided, by leaving a section of data out of the training process as testing data. Doing so is called cross-validation. To not produce any bias based on which data is selected for this testing process, k-fold cross-validation can be applied. This process consists of (randomly) splitting the data in k sub-sets, and using all but one fold to train the model and use the final fold to evaluate the performance measure. The result is then the average that is achieved for doing this process for each fold.

When comparing multiple models or tuning parameters, for example implementing different levels of regularization, this process can be nested to two-layer cross-validation. The inner set will split once more in k_2 folds. This splits the data into a training set and a validation set. The validation set is used to evaluate and choose the optimal model. The test set is then used to evaluate how well this optimized model performs on an unseen data-set. This means that k times a model is optimized and selected, where all of them will get an individual expected generalization error. This can provide a basis for choosing a model to use. A special variant of this k-fold cross-validation is leave-one-out cross-validation where k is set to the data size.

1.2 Position of the thesis in current framework

For purposes of animating plant based structures, Diener, Reveret, and Fiume, 2006 has applied image analysis through feature tracking for small plants indoor with a homogeneous background. In this case, the plant could be removed to allow for background subtraction. It was found that applying this technique to an outdoor tree proved difficult as features were rapidly being occluded and reappearing leading this method not being suitable for outdoor trees.

Further, an investigation towards a field with alfalfa wheat has been performed by means of particle image velocimetry by Py, Emmanuel De Langre, et al., 2005. Applying such a method to the tree however, assumes small displacements of particles within a grid cell. This would mean that either the resolution would have to be limited, which could cause trouble with incoherent and turbulent motion of the tree, or high displacements could not be captured. Both of these would limit the research of movement during high wind speeds. This could be mediated with a higher temporal resolution however.

A further investigation towards the frequency response of a single alfalfa wheat has been performed by Py, De Langre E., and Bruno Moulia, 2006. Here a dark homogeneous background was created to create high contrast between it and the object of interest. This is similar to image analysis performed by Doaré, B. Moulia, and De Langre, 2004 and Farquhar, Wood, and Van Beem, 2000. The latter of which has however also stated: "The pixel coordinates were found using semi-manual techniques because a skilled human operator, (JZW) was more effective than pattern recognition algorithms, with regards to compensating for varying marker illumination, shadowing, and reflectivity".

Two cases of applying image analysis to better understand wind-tree interaction with a focus on investigating the proportionality between wind speed and drag force is made by Vollsinger et al., 2005 and Rudnicki, S. J. Mitchell, and Novak, 2004. In these studies the crown streamlining for drag relationships have been investigated for hardwood and conifer species respectively. These investigations have occurred in a wind tunnel under laminar conditions and not in an outdoor actual wind environment.

This thesis explores to which extent image aspects can be extracted from an immovable tree without the need to provide a homogeneous background. It does so by proposing and applying a systematic method to analyse the images and then find concurrence with expected behaviour. The concurrence that is sought after includes an increase in enclosed and tree surface area as the tree grows leaves as well as a decrease in porosity. Beyond this, concurrence is sought after whether the swaying of the tree can be captured by checking that its direction is dependent on wind direction and its magnitude is dependant on wind speed. Finally, the surface area is expected to be less dependent on the wind speed when the tree has fewer leaves.

Beyond finding this concurrence, it further tries to shed light on the influence of frontal surface area reduction on the perceived Vogel exponent by investigating the drag coefficient calculated with a stationary and a dynamic surface area. Furthermore, it attempts to explore possible additional relations between the image aspects that are extracted and wind statistics that have been measured. Both a broad time period of 6-months using 10-minute averaged statistics and a more specific 10-day period using 1-minute averaged statistics is analysed.

The study applies machine learning in the filtering procedures, with a main focus on the rain interference filter.

1.3 Experimental setup

For the needs of this study the data acquired during a 6-month period using 10-minute averaged statistics and a 10-day period using 1-minute averaged statistics is analysed. This analysis pertains to a tree located in Risø campus of DTU, located at 55.688695 latitude and 12.095874 longitude. For this analysis, wind statistics, strain gauge measurements, camera images and atmospheric conditions are used. A description of these statistics, how they were gathered and a general layout of the terrain is given.

1.3.1 Terrain

The site at which the investigated tree is located is depicted in Figure 1.11.

Figure 1.11: Overview of the large-scale terrain around the tree.

Figure 1.12: Overview of the site with the tree, sonic anemometers and camera indicated.

The tree is located at the eastern side of the Roskilde Fjord. East and north-north-west of the tree, some buildings are present. Most western directions however are covered by water.

A close up of the terrain around the tree, also indicating the location of the used sonic anemometer around it and the camera that has been set-up is shown in Figure 1.12.

The landscape around the tree moving from the western meteorological mast towards the eastern meteorological mast is slightly uphill. This elevation difference, a description of the surface roughness and a general overview of the terrain is described by E. Dellwik et al., 2019.

Regarding the positioning of the sonic anemometers, this is also described by E. Dellwik et al., 2019 as: *"Two twelve meter tall lattice masts with a square cross section of 0.3 m were placed upwind (15m) and downwind of the tree (10 m). The two masts were instrumented with an ultrasonic anemometer each (uSonic-3Basic, Metek Gmbh, Hamburg) at four meters height. The anemometers were placed on booms extending 1.6 m from the side of the mast in the 195o direction."* Here upwind considers the western direction, and downwind the eastern direction.

The tree that is researched is an open-grown European Oak tree or Quercus Robus L. This tree is deciduous, allowing for a comparison between full-leaf and low-leaf conditions. It's height is around 6 meters and its topography is better described by Angelou, Ebba Dellwik, and Mann, 2019 where a 3D point cloud has been made to map the entire tree.

The tree is expected to grow to it's full-leaf condition around the end of May, however, around this period, the tree suffered an infestation by bugs, halting it's growth momentarily. Some of these bugs and tree damage is seen in Figure 1.13.

Figure 1.13: Pictures taken around the beginning of June 2019, showing bugs on the tree.

Regarding the camera used, the model is a DS-2DE4A220IW-DE, Hikvision with a resolution of 1920x1080. It's operational temperature ranges from -30 to 65 degrees

Celsius and it's operational humidity is up to 90%. It includes a night mode with longer exposure time that produces gray-scale images.

1.3.2 Wind resource

Wind resources are gathered by the previously described sonic anemometers. These have a sampling frequency of 20 Hz and save their statistics in either 1-minute or 10-minute averages. Both the mean and the standard deviation of the wind speed is recorded. Two time spans are investigated, both will be discussed separately.

For each meteorological mast, for some direction interval, it will be in the wake of the tree. To get the inflow wind speed, a switch statement is used to determine from which meteorological mast the data is retrieved. As the meteorological masts diametrically oppose each other around the tree at 110 and 290 degrees, data from the meteorological mast angle closest to the wind direction is chosen. The wind direction that determines the switch statement the western meteorological mast is used. In other words, when the wind direction comes from the direction range between 20 and 200 degrees, the eastern meteorological mast is used and for other cases the western meteorological mast is used.

It is noted that after the installation of the sonic anemometers, a misalignment of 6 degrees was discovered. This misalignment has been corrected.

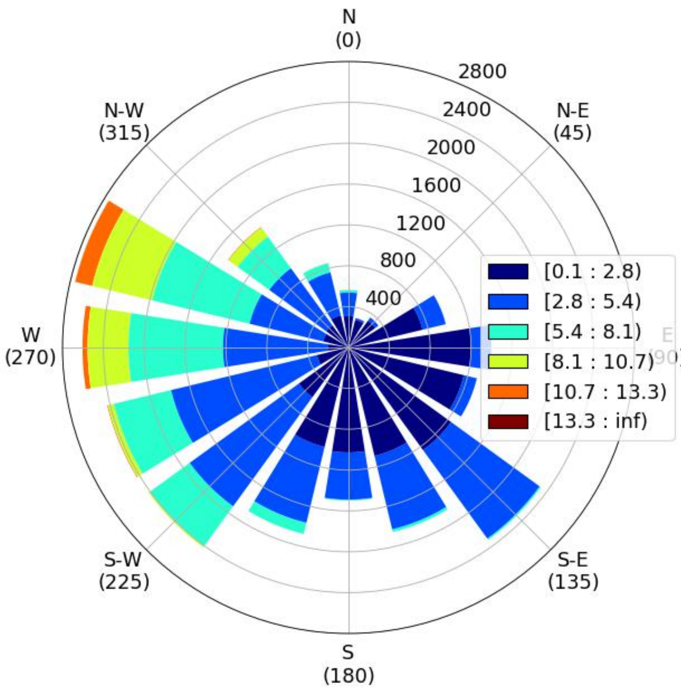

Figure 1.14: Wind rose showing how many 10 minute averages are available binned in wind direction and wind speed.

6-Month time span

The 6-month time span looks at 10-minute averages and ranges from the first of May 2019 until the 31-st of October 2019. This time period has a gap in data between 12:30 on October 18th until 11:00 on October 27th. The wind rose belonging to the wind statistics of this time period is shown in Figure 1.14.

Mostly western or south western wind speeds are seen. Western wind speeds see a higher velocity profile than eastern wind speeds, most likely due to the surface roughness associated with western directions (internal boundary layer between water and grass) vs the surface roughness associated with eastern or northern directions (high shrubbery or buildings). The distribution of wind speeds together with a Weibull fit is seen in Figure 1.15.

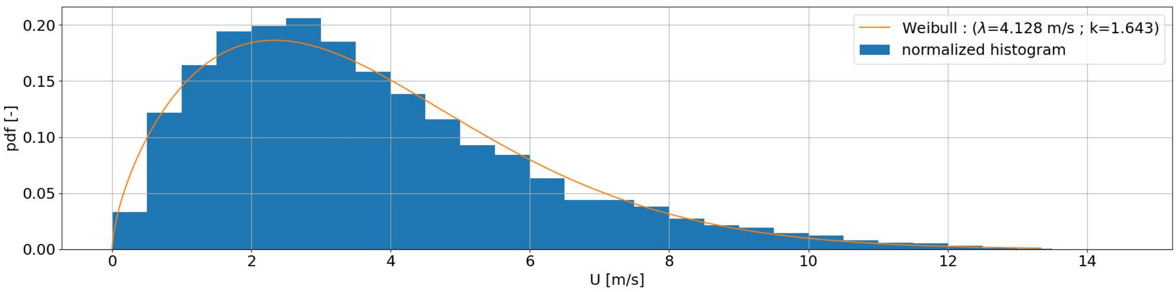

Figure 1.15: Normalized histogram of the wind speeds with a Weibull distribution fitted to it.

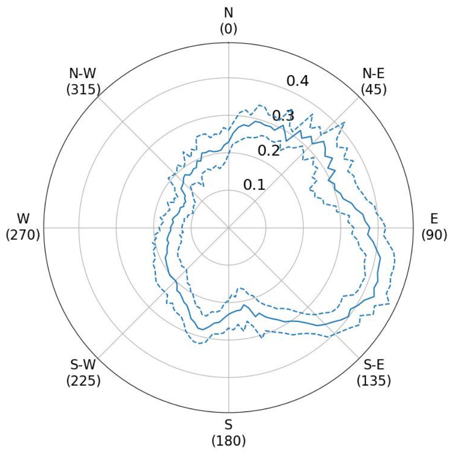

Figure 1.16: Polar plot showing the turbulence intensity per wind direction. The dashed lines represent one standard deviation above and below the mean.

A somewhat standard wind profile is seen which is fit by the Weibull well. It is seen that wind speeds range up to 14 m/s, and that in the high speed regime, as speeds get higher, less data is available. As for the turbulence intensity of the wind, a polar plot of this is seen in Figure 1.16 where wind speeds are binned in 128 equidistant bins. It is noted that wind speeds below 2 m/s are not included in this, as they inhabit dis-representative high fluctuation in turbulence intensity.

For wind speeds coming from the west, much less turbulence intensity is seen, which is again attributed to the directional difference in surface roughness.

10-Day time span

During the 10-day analysis, a specific wind sector is investigated where the surface area that is recorded by the tree is considered the frontal surface area. For this, the western meteorological mast will always be used for gathering wind statistics. At some points however, the wind deficit will be investigated at which point the data from the second mast will also be used. Wind statistics posed here, are those of the western meteorological mast. The wind rose for the 10-day analysis is seen in Figure 1.17.

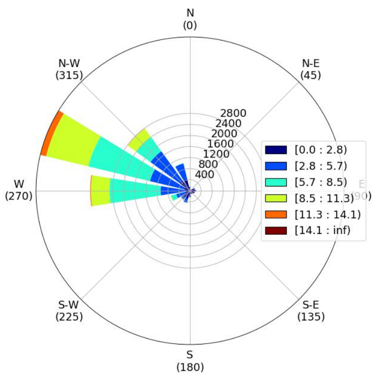

Figure 1.17: Wind rose showing how many 1 minute averages are available, binned in wind direction and wind speed.

Comparing this figure to Figure 1.14, much more westerly directions are seen. This is no coincidence, as the 10-days for inspection have specifically been chosen due to the high occurrence of westerly winds across a large velocity interval.

The wind speeds available and the fraction of these wind speeds that lies in the western section centered around 290 degrees is seen in Figure 1.18. Note that no Weibull density function is fitted to this data-set as the Weibull distribution is mainly a good shape for a distribution of 10-minute mean wind speeds.

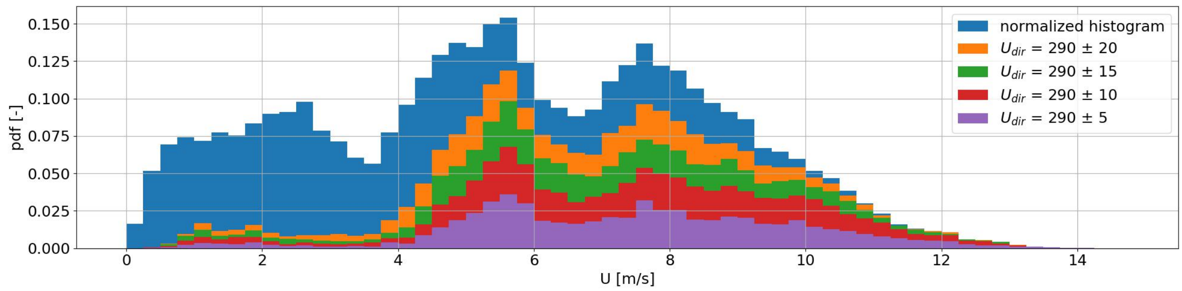

Figure 1.18: Normalized histogram of the wind speeds. Also showing fraction of wind speeds centered around 290 with margin of 20, 15, 10 and 5.

It is seen now that while many data is present in the higher speed ranges, this is not the case for speeds below 4 m/s when filtering out any non-westerly winds. It is also seen that increasing the margin from 10 degrees to 15, does not improve the number of high wind speed data points by a large margin.

1.3.3 Atmospheric conditions

For purposes of the 10-day analysis, the density is also considered. The density signal throughout the 10-day period is seen in Figure 1.19. It is seen that throughout a day, fluctuations up to 3% can occur.

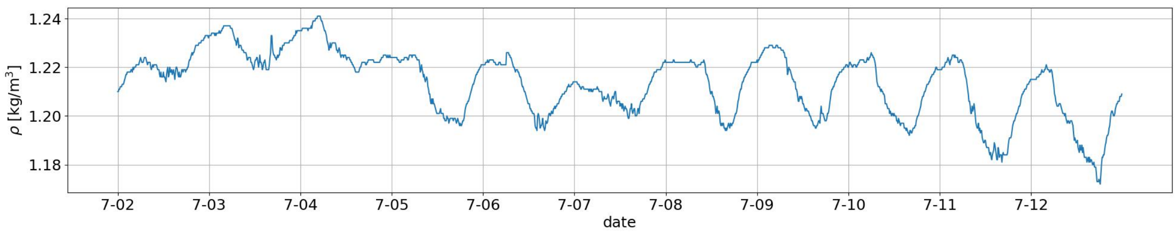

Figure 1.19: Time series of the density (ρ) signal for the 10-day period.

1.3.4 Force measurements

As part of a previous study (Angelou, Ebba Dellwik, and Mann, 2019) a strain gauge has been set-up on the tree, and by pulling the tree with a known force on a non-windy day, this strain-gauge has been calibrated. Due to temperature and diurnal effects on and of the tree, this calibration will drift throughout the day. Furthermore due to the growth of the tree, this calibration will also drift throughout it's growth. Correction factors for this have been researched and created by Angelou, Ebba Dellwik, and Mann, 2019 and the strain data for the 10-day period is transferred to a bending moment on the tree.

Then assuming that the center of the force acting on the canopy is at a height of 3.9m, the arm between this point and the location of the strain gauge can be determined. This allows for the calculation of the drag force acting on the tree. The force signal for the 10-day period can be seen in Figure 1.20.

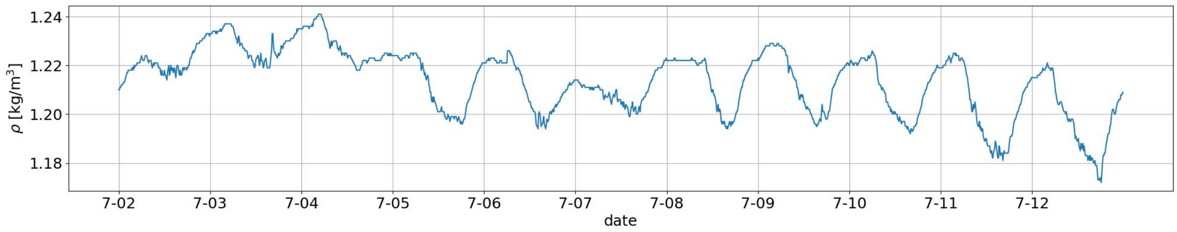

Figure 1.20: Time series of the drag (F_D) signal for the 10-day period.

CHAPTER 2
Methodology

The ultimate goal of the methodology is to create tools/methods that allow for the comparison of image statistics with wind statistics. All methods developed and/or applied for this comparison are described in this section. These methods are split up in three main sections. Methods that aid in locating of the tree, methods that allow for the extraction of image aspects and methods that filter out unusable images. The results of these methods and the linkage between image aspects and wind statistics will be described in the results section of this thesis.

2.1 Tree masking

Besides the tree, the image also includes the sky with clouds, ground with grass and water, and even the eastern meteorological mast which partly occludes the tree. The object of interest however is the canopy of the tree. To ensure that these other factors do not interfere with any image processing, a preliminary mask focusing on the canopy is set up. This preliminary masking consists of three steps which are seen in Figure 2.1. These steps are: cutting off the horizon, masking out the foreground objects (meteorological mast and at some point a pole which is positioned at the bottom left part of the tree) and drawing an ellipse around the canopy.

Figure 2.1: Depiction of step-wise process of masking the tree. First masking out the horizon, then masking out the meteorological mast and finally masking out the sky.

During the two year period where images have been made of the tree, repositioning of the camera has occurred. Furthermore, the tree is deciduous making it resize throughout the seasons to the point where no single mask can be used to locate the tree. This repositioning of the camera is seen in Figure 2.2. From this it becomes evident that a different mask is needed after each repositioning of the camera or when the tree grows beyond the current mask. To do this manually, would be tedious and laborious. This suggests the need for an automated way of applying the three masks shown in Figure 2.1.

Figure 2.2: Two pictures made of the tree at two different times, indicating the repositioning of the camera.

As the movement of the camera on a daily basis is minimal, this process will be done for each day. Setting the mask will require a reference image. If any single snapshot would be used to set the mask for a day, this could be at a moment where the tree deviates from its standard position due to a gust or something interfering with the image, leading to a sub-optimal mask. For this reason and filtering purposes which will be discussed later, a mean image is created.

This automated method will be made to work for days where the image shifts less than 10 pixels. This means that a day will be omitted from analysis if any reference point that aught to be stationary moves more than 10 pixels during a day. For stationary reference points the foreground sonic anemometers are used.

2.1.1 Mean image

As has been stated, a mean image will be created for this analysis. A mean image is created by firstly eliminating all nighttime images and images where data is missing. From the remaining images, the average of the blue channel is taken as this channel has been seen to provide the best contrast between the tree and the sky.

However, creating a mean by using all images in a given day, would be computationally expensive. A comparison of the IOU (intersection over union) and computation time (t) of lower fidelity mean images with a high fidelity mean image is performed. As a high fidelity image, a binarized image of the mean of all images in a day is used. For low fidelity mean images, every N^{th} image is used to create the mean image. To decrease uncertainty in computation time, the process is ran 5 times and the mean computational time is depicted. The result of this investigation for values of $N = 1, 4, 16, 64, 256$ is shown in Figure 2.3. From the results of this analysis, $N = 16$ is chosen as an optimal value. This is as it is seen that using every 16^{th} image leads to an IOU higher than 0.99, while cutting computation time by a factor of roughly 16.

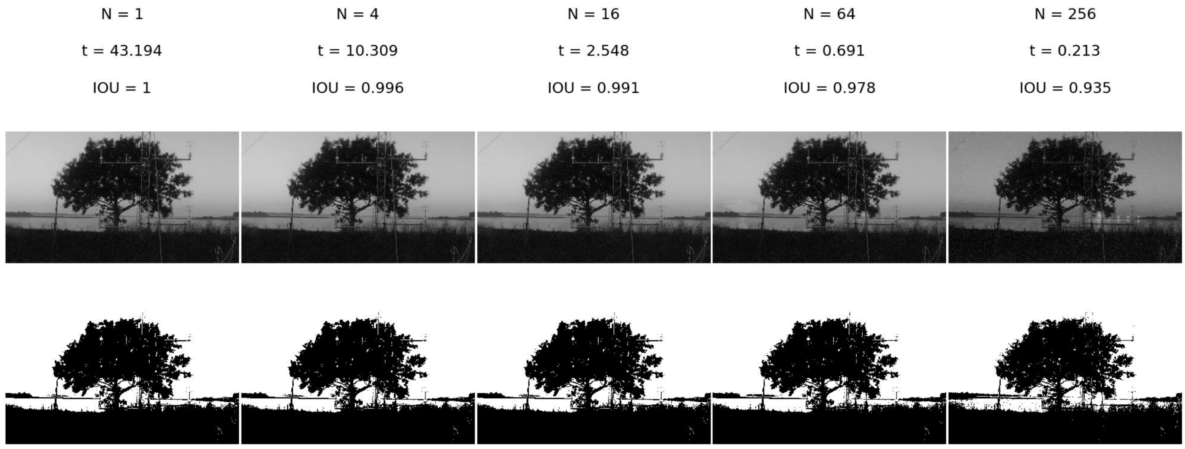

Figure 2.3: Comparison of computation time and IOU by using every N^{th} image for creating a mean image. The top row of images represents the mean image, and the bottom row represents the result of binarizing these images.

2.1.2 Masking below the horizon

A split is made at the horizon to separate earth from sky. In order to find the horizon of the image, the left most 20 columns of pixels of the mean image are investigated. Nest, a the first spike in a change measure is associated with the horizon. This change measure looks at the difference between n values with their n values n steps further. Which means that this change measure spikes when a big shift in mean value is found that persists over n steps. The first occurrence of this when approaching from the top side, should be the land mass seen on the left of the image. This measure y_{change} is shown in Equation (2.1). In this equation, f represents the vertical line of mean values.

$$y_{change}(x, n) = \left| \sum_{i=x}^{x+n} f(i) - f(i+n) \right| \tag{2.1}$$

The result of this method for a mean image is seen in Figure 2.4. The limit is set twenty pixels lower to take into account that the landmass seen on the left side of the image is slightly higher than the horizon. Everything below the horizon line is masked out as the first part of the preliminary mask.

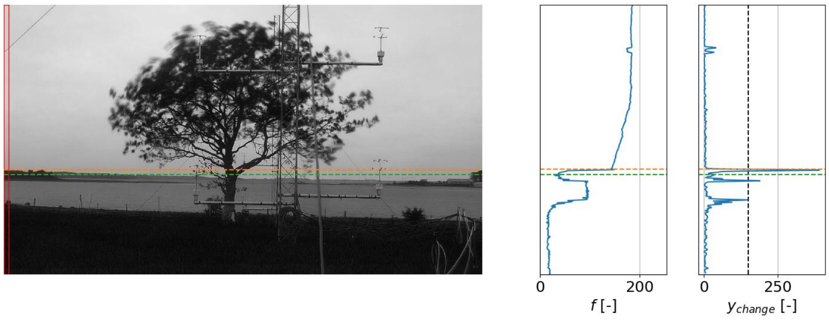

Figure 2.4: Localization of the horizon by investigation of the change measure (right) of the mean of the left most 20 pixel columns (middle) indicated with the red rectangle. The orange line indicates the first time the change measure exceeds the limit, and the green indicates the margin that is added.

2.1.3 Masking the foreground

Besides the earth below the horizon, the foreground objects will also interfere with image extraction and thus need to be filtered out. To do this, the pixels belonging to the foreground objects need to be identified. For this purpose, two hand drawn masks are created on a reference image. These two masks can be seen in Figure 2.5.

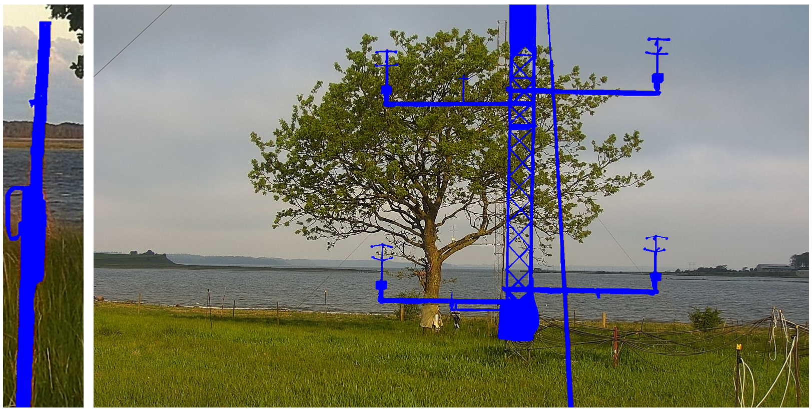

Figure 2.5: The hand drawn mask indicated in blue for the foreground pole (left) and meteorological mast (right).

For both of these masks two positions are marked. For the pole these are at the top and in the middle and for the meteorological mast these are at the bottom left and top right sonic anemometer. Next, the pixel coordinate of these objects in the mean image for each relevant day is noted down. The mask is then resized and repositioned in such a way that the marked locations in the masks coincide with the pixel coordinates that have been noted down.

In this process, two observations are made. Firstly, if the meteorological mast mask is shifted down, this means that the mask will shift below its top markings, in this case anything above the top most available point of the mask is masked out as well (indicated in the right side of Figure 2.6). Secondly, it is assumed that this method does not produce perfect results as the projection of the meteorological mast might slightly morph rather than only shift and scale. For this purpose, a margin of three pixels is added to the meteorological mast mask, and six pixels for the pole mask. An example of two images with the masks fit onto them is shown in Figure 2.6.

Figure 2.6: Two examples of removing the horizon and the foreground.

An investigation is performed to see if no camera shift occurs throughout the day such that the meteorological mast escapes the mask. From this investigation it is concluded that a significant shift does occur as the meteorological mast escapes the mask at some time throughout the day. A zoomed in view of the meteorological mast mask on the top right sonic anemometer is shown at three times in Figure 2.7. In this figure it is seen on the left most side that the mask does not cover the meteorological mast fully at some point as the camera has shifted slightly. In order to amend this, firstly, the slim version of the mask (where no 3 pixel off-set is added) is taken. Consequently, this mask is cross-correlated with the inverse of the gray-scale of the shown image. This means that a high count of dark pixels within a shifted mask (dark pixels belong to the meteorological mast) will correspond to a high value for the cross-correlation. The shift that maximizes this cross-correlation is applied to the slim mask and the mask is given the three pixel off-set again. This process is repeated on an hourly basis, as the mask is assumed to not shift more than three pixels within this time-frame.

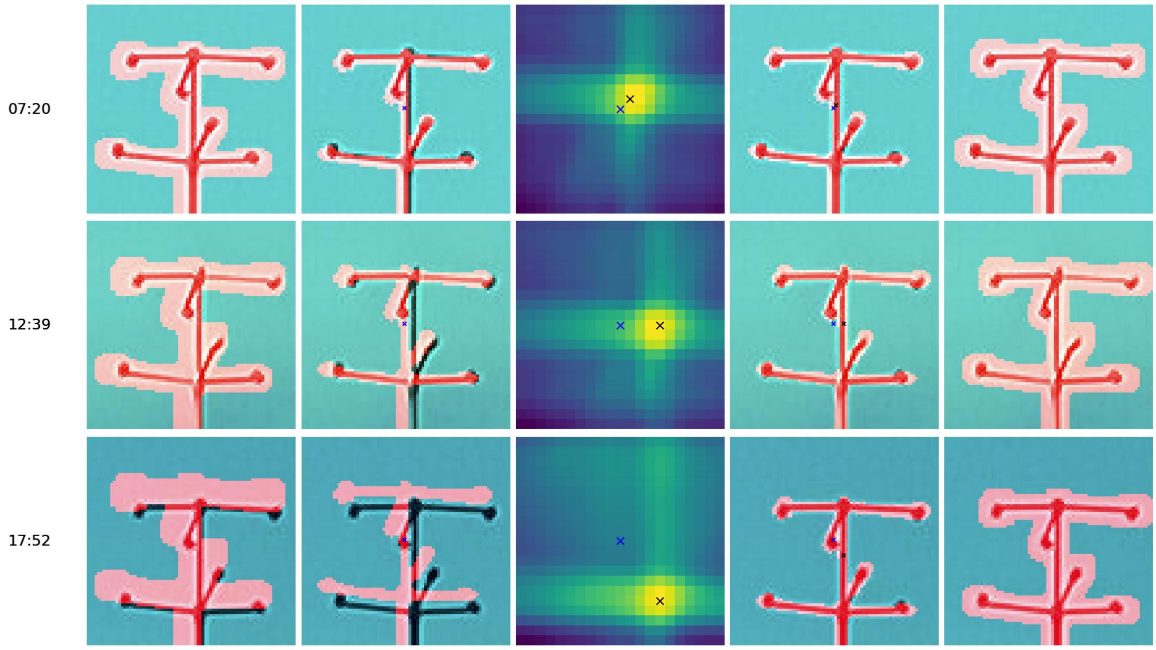

Figure 2.7: A zoom of the top right sonic anemometer for three time-stamps during a day where the meteorological mast mask is indicated in red. Showing the initial mask, slim mask, 2D-correlation matrix, corrected slim mask and updated shifted mask from left to right.

2.1.4 Masking the sky

The final part of the mask is an ellipse that marks the canopy and excludes a large part of the sky, where no information about the tree movement can be gathered. This ellipse is centered at the center of the canopy. To find the center of the canopy, a somewhat suitable solution would be to binarize the already masked part of the image, followed by finding the weighted average of the active class (tree). However, the tree is more dense on the left side, leading to the center being marked more to the left. To account for this, the tree is leniently down-scaled by a factor of 10 with a threshold of 0.03, increasing the weight of the less dense right section of the tree. A qualitative example of this is seen in Figure 2.8. If a factor lower than 10 is chosen, the purpose of leniently down-scaling is not fulfilled as the gaps on the right will remain as such. Besides this, the factor is preferred to be low to limit an introduced discretization error. The threshold is set to 0.03 to be lenient, but not lower, as otherwise the diagonal cable at the top left would introduce a noticeable bias on some days.

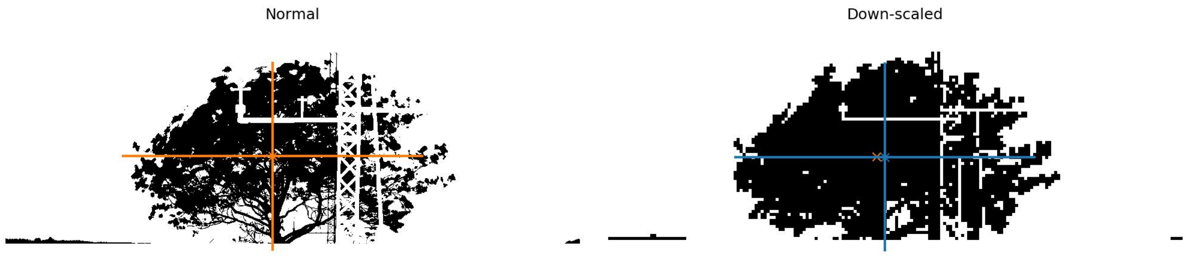

Figure 2.8: Visualisation of the difference in center for the normal (left) and down-scaled (right) case.

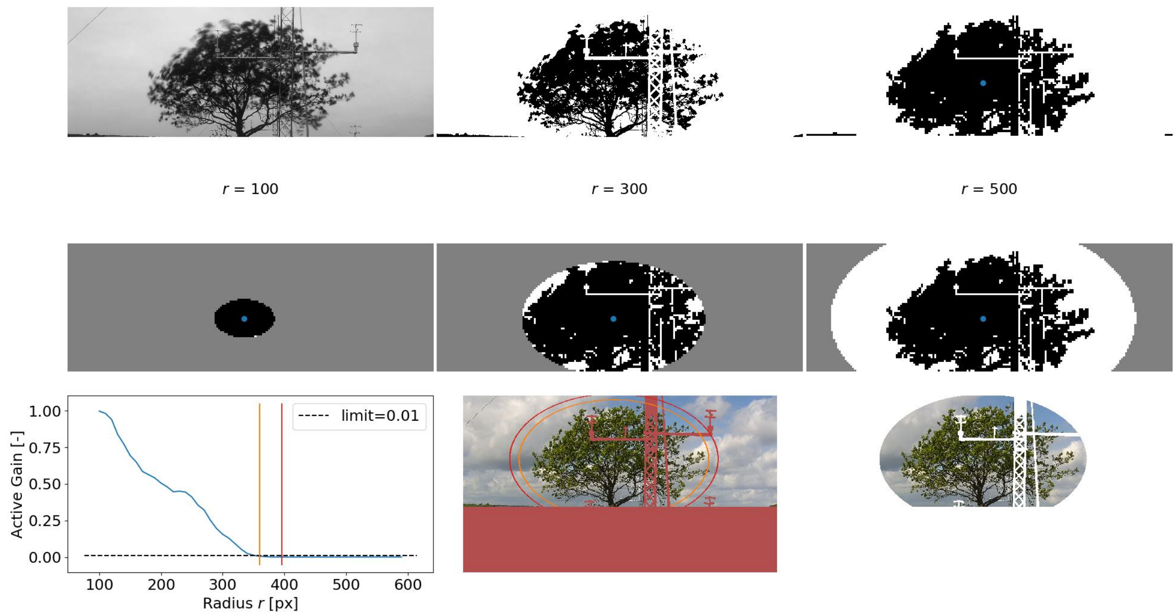

Figure 2.9: Visualisation of the process of finding the preliminary mask sequentially from left to right, top to bottom: 1.) cut off horizon 2.) binarize image 3.) downscalee the image 4.) visual of $r = 100$ 5.) visual of $r = 300$ 6.) visual of $r = 500$ 7.) plot to find limit intersection 8.) add margin 9.) final mask.

This allows for the finding of the location of the center of the ellipse. Next, the aspect ratio and radius need to be determined. The aspect ratio is set at 1.6 for this specific tree based on visual inspection. The radius is determined by drawing an increasingly bigger ellipse (increasing the radius by 10 every step) and noting the fraction of the added pixels that are active over total added pixels by increasing this radius. As soon as this fraction is less than a limit (set at 0.01) it can be said that increasing the radius more, does not significantly add more parts of the tree. To add a margin for movement of the canopy and camera shifts, the radius is increased by a factor of 1.1. This whole process is shown in Figure 2.9.

This process is validated on multiple images and is seen to accurately determine the location of the canopy. An example for two different situations is shown in Figure 2.10.

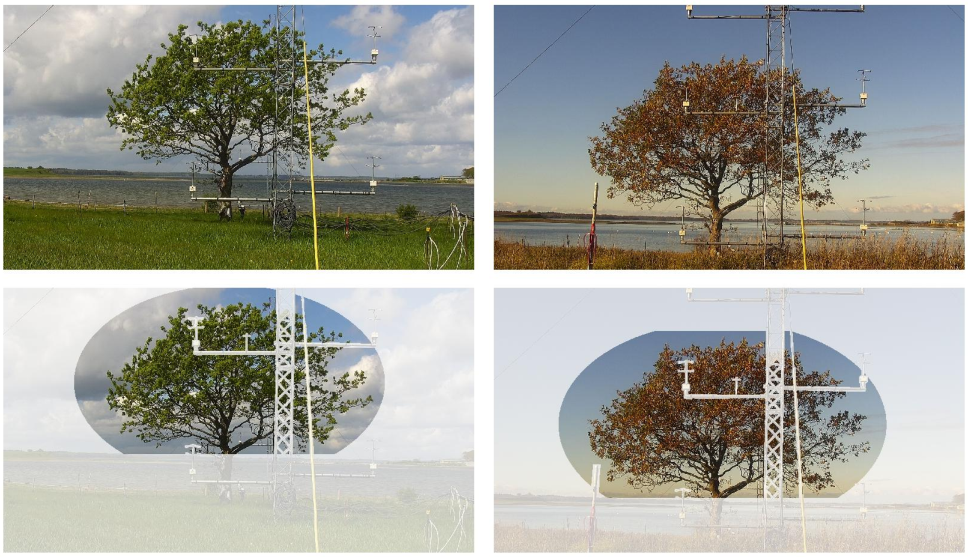

Figure 2.10: Two examples of full masking of the tree.

2.2 Image aspect extraction

Four image aspects are to be extracted. The tree surface area, the enclosed surface area, the horizontal center of area and the vertical center of area. By extracting both the tree and enclosed surface area, by proxy, the optical porosity is extracted.

These aspects are extracted by binarizing the tree and finding the circumference of the binary image. Then for the enclosed surface area, the area within the circumference is determined. For the tree surface area, the pixels belonging to the tree within the enclosed surface area are counted, in addition to the canopy density extrapolated onto the meteorological mast. These surface areas need to be converted into a physical area by knowing the surface area that one pixel represents. This is found by estimating the distance between the tree and the camera. With this the physical location and size of each tree pixel is known, also allowing for finding the center of area. By proxy, the optical porosity is also found when retrieving the tree and enclosed surface area.

2.2.1 Binarization

The goal of binarization is to separate pixels into two classes. Pixels belonging to the tree, and pixels not belonging to the tree. This process is done by means of thresholding which is made automatic by setting the threshold at the point of maximum in-between class variance as described by N., 1979. The channel used for this thresholding was recommended to be the blue channel by Lee, Alfaro, and Van Sickle, 1983 as it provided the most contrast between the tree and the sky. However, the possibility of a better channel combination is investigated.

The main requirement for thresholding is that the 3D image (x,y,color) is transformed into a 2D image by taking a combination of each color channel. This transformation must occur in such a way that thresholding the image leads to the best split between tree and non-tree pixels. To do this quantitatively, three sections of the tree have been labeled to create a ground-truth. These sections and their labels are shown in Figure 2.11. The three images that have been chosen are that of a daytime tree in the shadow of a cloud, a daytime tree in the sun, and an evening sun. The regions that have been chosen are chosen at the locations where simply using the blue channel was found to perform poorly at correctly thresholding the image.

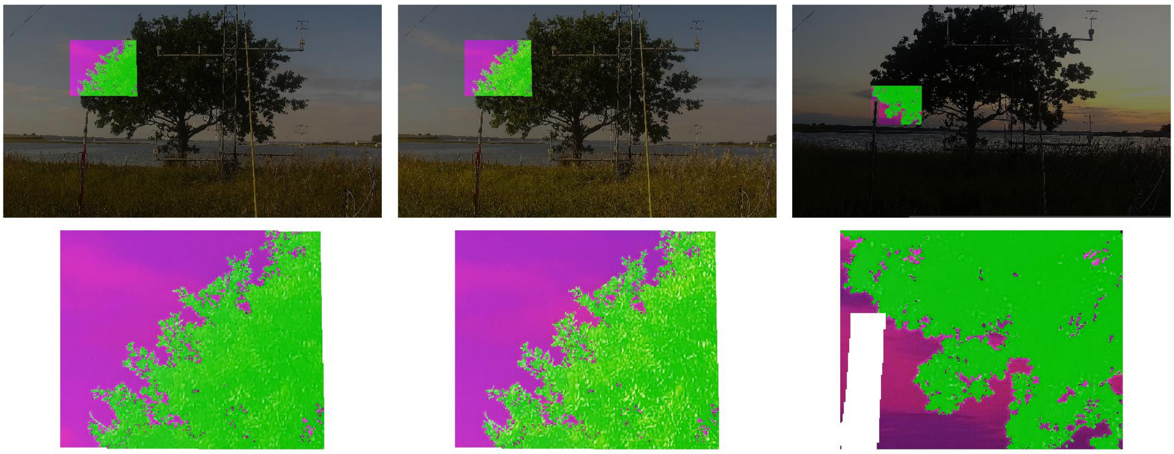

Figure 2.11: Three image labels used for selecting the best channel combination Each image represents a ground truth to be compared with. Greenish regions indicate pixels that should be classified positive, purplish regions indicate the opposite.

To find the best channel combination, the blue channel is set to one, and the amount of red and green channel that is added or subtracted is set as a variable. These two variables are then explored with the optimization algorithm described by Nelder and Mead, 1965. The optimization problem is to minimize the false negative plus the false positive labels of all images. This optimization algorithm is chosen as the performance parameter is discreet and perhaps noisy which disallows the ability of discreetly finding the gradient through tactics such as finite differences when the step is too small. The results of this exploration and comparison between the blue channel and the optimum is seen in Figure 2.12. It can be seen that whereas using only the blue channel interprets the high-value reflections in the tree of the illuminated image (middle) as gaps in the tree, the optimized image does so to a lesser extent.

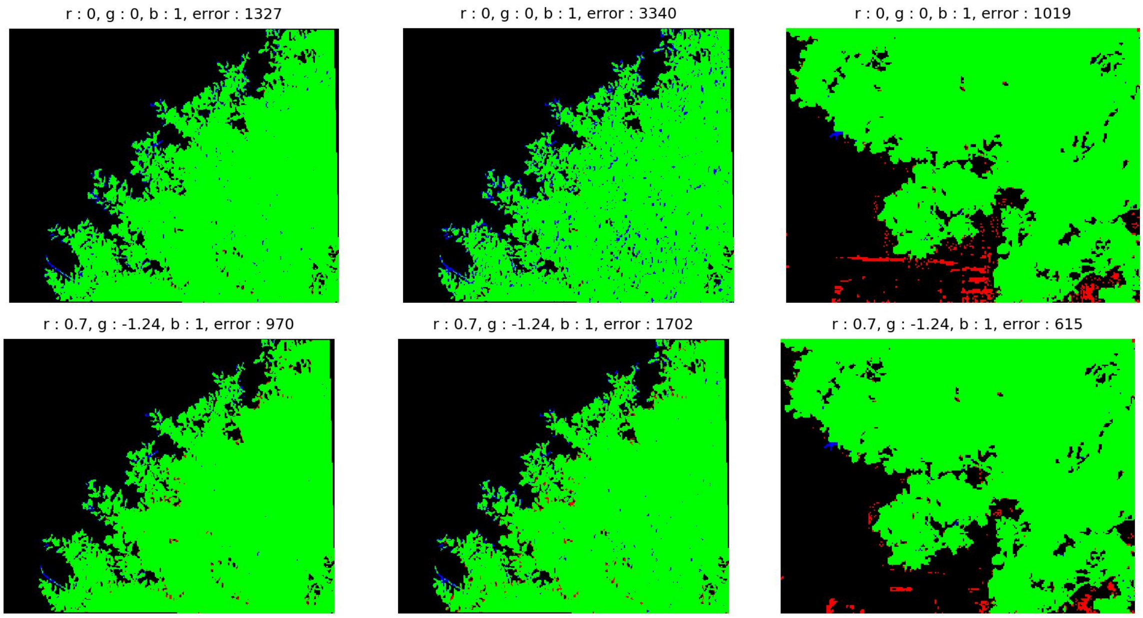

Figure 2.12: Comparison between the optimal channel combination and using the blue channel for thresholding (true positive: green, false positive: red, true negative: black, false negative: blue).

This means the tree is binarized using automatic thresholding on a channel combination of 0.7 times the red channel, -1.24 times the green channel and 1 times the blue channel.

2.2.2 Finding circumference

The circumference is required to determine the enclosed area. The circumference is found by identifying all pixels that enclose a large shape in the binary image. These pixels are found by putting an agent on the binarized image and letting it walk around the positive class until it reaches its starting point once more. It will then return the coordinates of the points if the number of them is above a threshold. If it is below a threshold, it means that the shape it encompassed is most likely not the mayor shape in the image and it will try again. This entire process is shown in an algorithm in Appendix A.

This algorithm is visually represented on a simple shape in Figure 2.13. Here the blue line indicates it's searching direction. The result on multiple arbitrary simple shapes is shown in Figure 2.14. Here the blue dots represent the hull points that have been found.

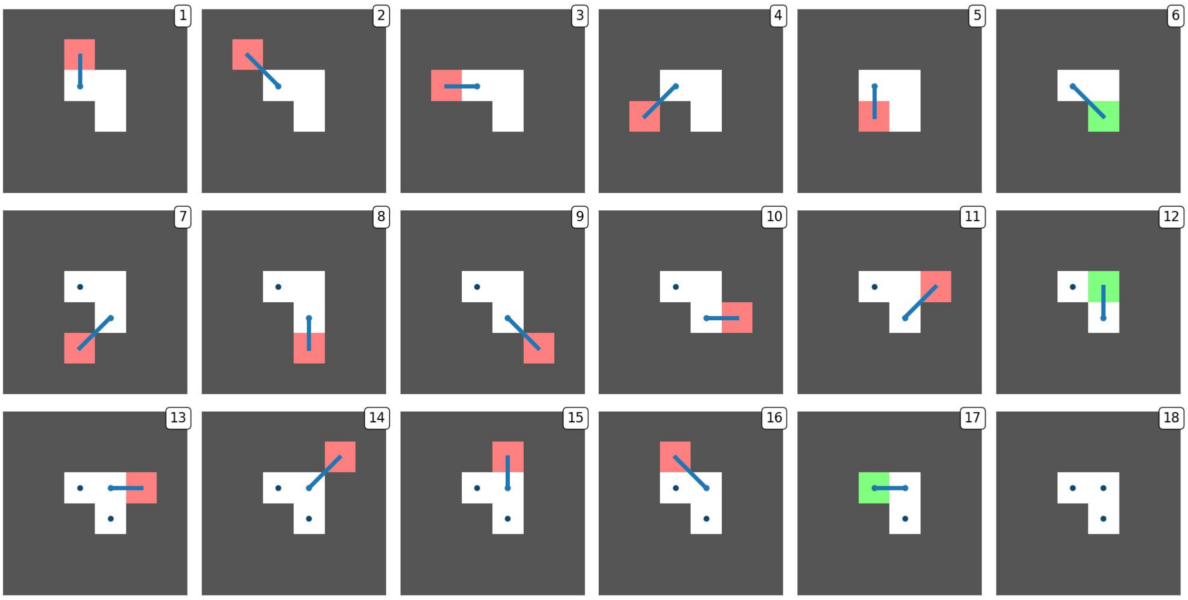

Figure 2.13: Circumference walking algorithm on a simple shape. The blue line indicates the link between the current point and the inspection point.

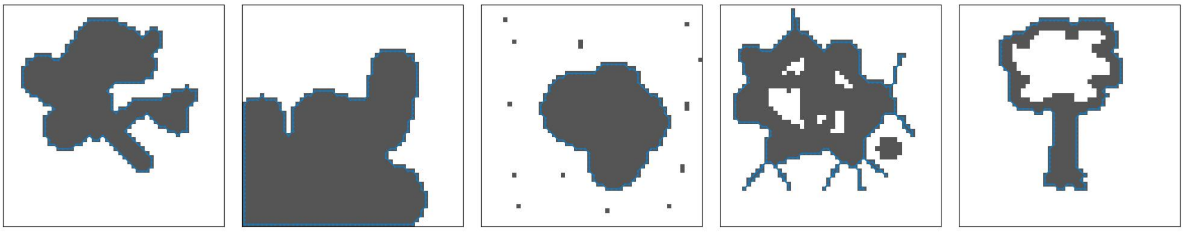

Figure 2.14: Examples of the circumference points (indicated in blue) found for five arbitrary shapes.

2.2.3 Enclosed surface area

The enclosed surface area is found by drawing line segments between each of the circumference points. Every pixel below each of non-vertical line segments will be switched. Meaning that if it is off, it will be turned on, and if it is on, it will be turned off. If the line segment is moving towards the left, it will include pixels on the line segment, and if it is moving to the right, it will exclude pixels on the line segment. This process is seen in Figure 2.15. When moving to the left, this process is exclusive of the first point, unless the previous segment was not moving to the left (step 15). Segments moving to the right are exclusive of the last point, unless the next segment does not move to the right (step 11).

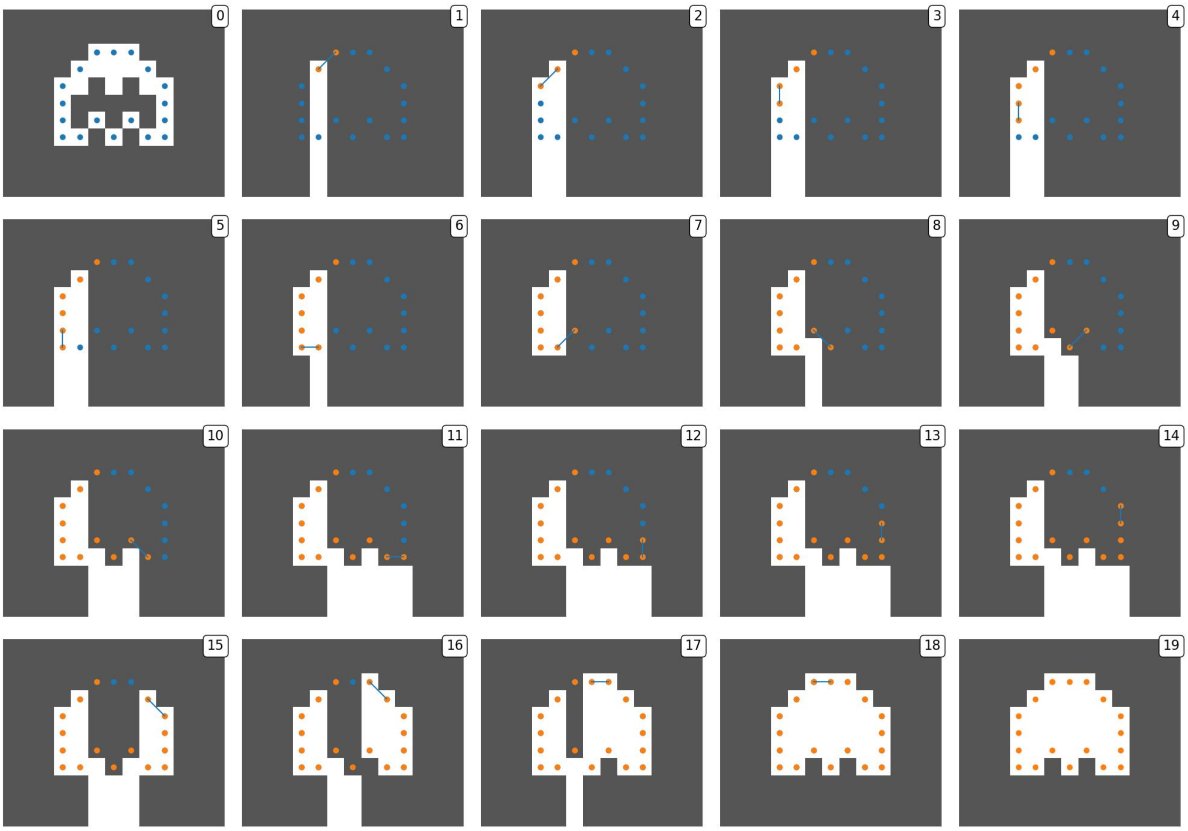

Figure 2.15: Step-wise method of finding the enclosed area given the circumference points.

This process can also be applied when points are spaced out more. This is seen in Figure 2.16. The same procedure as in the previous example is applied with one additional final step. The spacing of the points, leads to the enclosed area being more tightly wrapped around the circumference points. To more closely resemble the initial image, an offset equal to half the point spacing is applied. This is done by switching on any pixel which has a neighbouring active pixel (diagonal included) for an amount of times equal to the required offset. This is seen in the final step in Figure 2.16.

This process is applied to the binary image of the tree, however, an issue arises. If the regular binary image of the tree is used, the circumference walking agent will walk into the tree to a high extent as it is prone to pixel-wide gaps in the tree. Furthermore, some parts of the tree might be attached with thin branches, classified as sky. This will make parts of the tree be detached from the main shape by a little amount. A remedy for this, is to use the down-scaling described in the introduction with a lenient threshold (0.1). This will fill-in any small gaps that are in the tree and allow the agent to walk over them. This process for different factors of down-scaling r is seen in Figure 2.17.

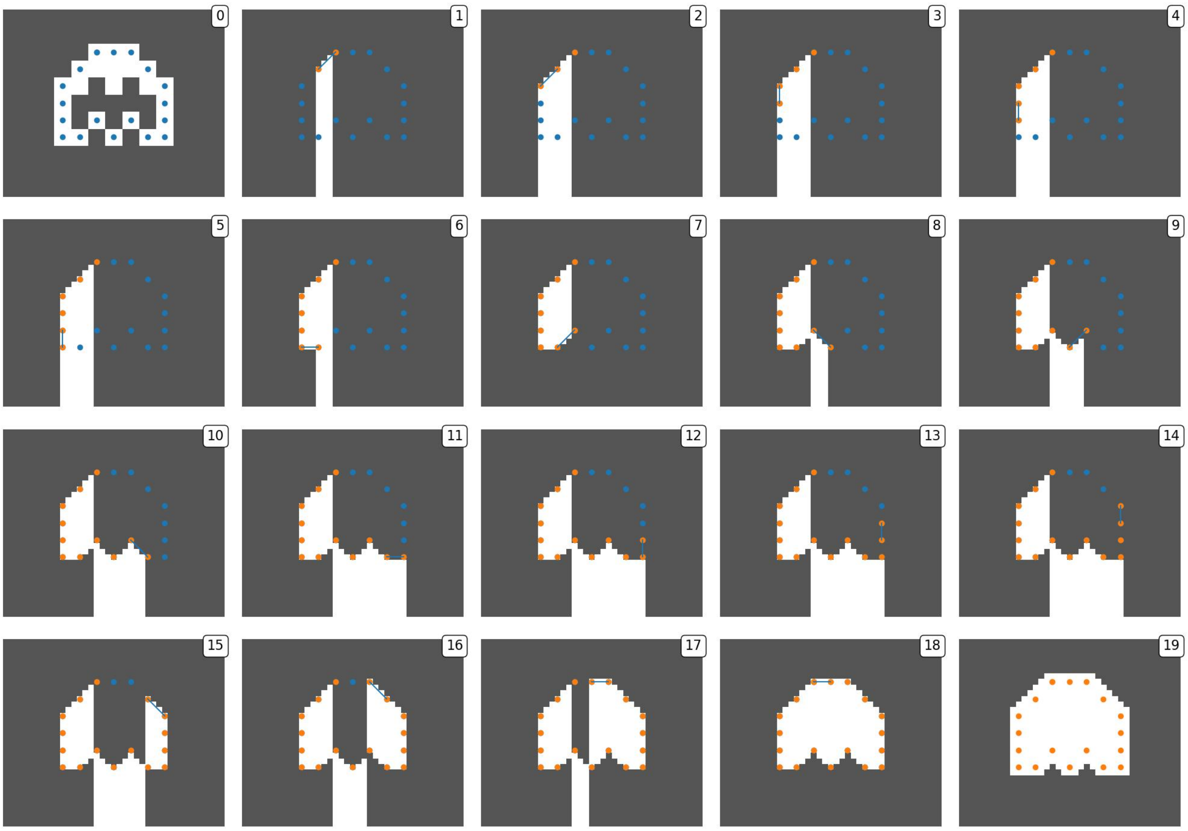

Figure 2.16: Step-wise method of finding the enclosed area given the circumference points for larger circumference point spacing.

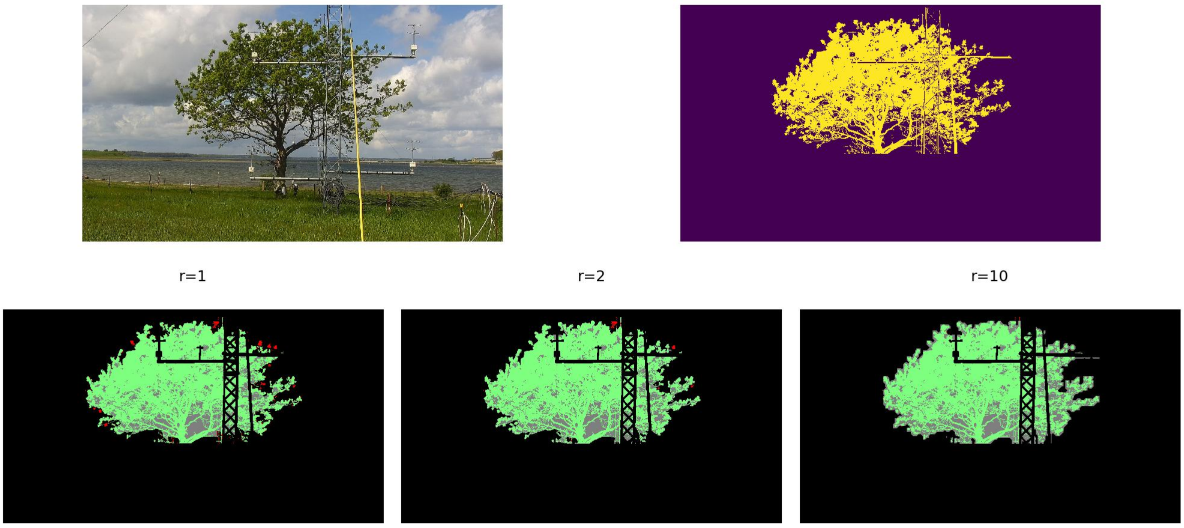

Figure 2.17: Comparison of different values of r in down-scaling, where red depicts disregarded areas, gray depicts the area mask and green depicts the tree within the area mask.

Although this remedies the agent from missing parts of the tree and not walking into the tree, the choice for r is arbitrary though it impacts the total surface area that is found. This is as gaps of all scales are present in the tree. No threshold can be clearly set, without creating some bias towards what to include as being part of the tree. The effect of this to an extreme case is seen when inspecting the lower extruding branch of the tree in windy conditions. The enclosed area for two images taken one minute apart, by using $r = 5$ and $r = 10$ is seen in Figure 2.18. Here also a zoom is given to the discussed branch. It is seen that changing the value of r has influence on the total surface by roughly one percent, which is significant when compared to the influence of the wind.

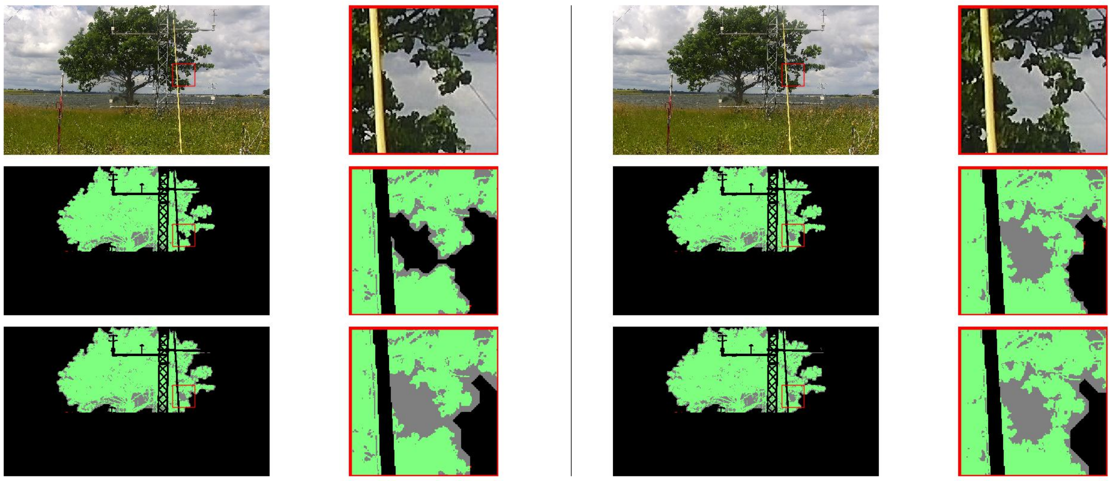

Figure 2.18: Enclosed area for two images taken one minute apart from each other shown in the left and right side. The top row shows the image, the second row shows the enclosed area by using $r = 5$ and the third row shows the enclosed area by using $r = 10$.

The question then becomes what non-tree parts of the image should be considered as being part of the enclosed area. For this thesis, a value of $r = 10$ is used to decrease the fluctuation seen in Figure 2.18. A serious note is made however, to the arbitrary nature of the enclosed surface area, and by extension thus also the optical porosity. All analysis regarding this feature should be relative or interpreted with reservations, as the absolute value is highly dependent on the length scale of r. The determination of the circumference of the tree is analogous to the coastline paradox described by Mandelbrot, 1967.

2.2.4 Tree surface area

Regarding the relation with wind statistics, both the enclosed and the tree surface area will be extracted. The true value of both however can not be known due to the occluding meteorological mast in front of the tree. So far, this area has been masked out and is regarded as being part of the sky as occlusion makes it impossible to state with certainty which pixels behind it are part of the tree. A guess can be made however. This is done by applying a five step process. This process is seen in Figure 2.19.

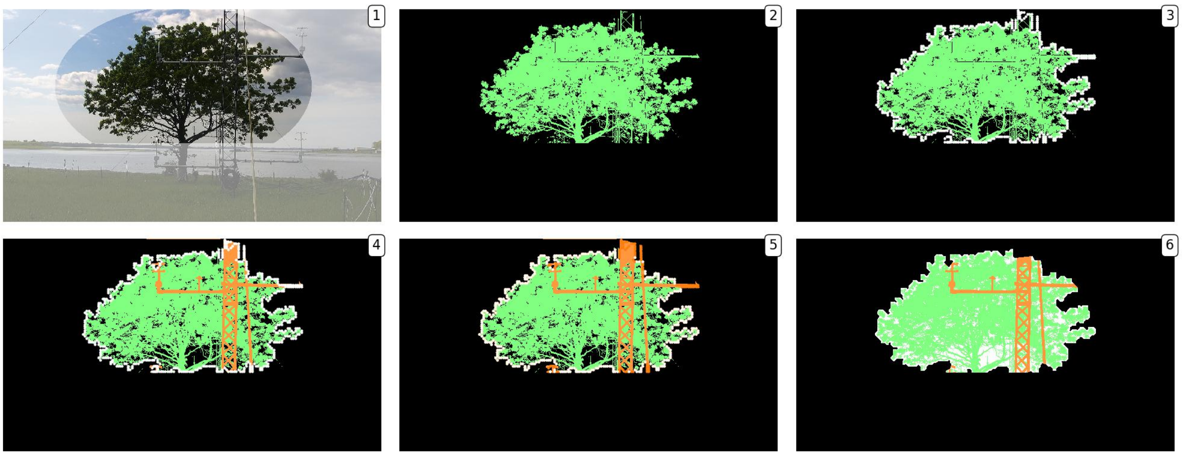

Figure 2.19: Six step-wise images. 1.) The normal image with horizon and sky masked out 2.) The binarized image 3.) The circumference points are found 4.) The meteorological mast is identified 5.) Circumference points within 5 pixels of the meteorological mast are excluded 6.) The enclosed area is found and pixels are classified.

First the image is binarized. Second, the circumference points are found as per the method described in Section 2.2.2. Next, the meteorological mast mask is applied. All circumference points within 5 pixels of this mask are removed. Following this, the enclosed area is created with the remaining circumference points, effectively cutting off the meteorological mast where it is expected to not have any parts of the tree behind it. Finally, the areas are identified as seen in the bottom right image (6) of Figure 2.19. The white area is described as background (S_w), the green area is described as tree (S_g) and the orange area is described as the meteorological mast (S_o). The final enclosed surface area S_e as it will be used in the results section, is described by the sum of all these areas. The final tree surface area S_t as it will be used in the results section, is described as the tree area outside the meteorological mast plus the estimated canopy density extrapolated onto the meteorological mast area. This is seen in Equation (2.2) and Equation (2.3).

$$S_e = S_g + S_w + S_o \tag{2.2}$$

$$S_t = S_g + S_o \cdot \frac{S_g}{S_w + S_g} \tag{2.3}$$

2.2.5 Pixel area conversion

When the tree is binarized, the area is extracted in the form of amount of pixels. This means however, that when the distance between the camera and the tree is changed, this will be noted as an increase or decrease in area. To find a more physical and continuous representation of the surface area, the number of pixels is converted to m^2. To do this, the surface area or length of one pixel is required.

To find the length of one pixel, first the length of the width of the screen at the depth of the tree is found. As there are 1920 pixels in this width, the length of one pixel would be $1/1920^{\text{th}}$ of this width. The width of the screen at the depth of the tree is found by looking at Figure 2.20. From this, the relationship in Equation (2.4) can be used to find the width when the horizontal field of view (γ) and the depth or distance between camera and tree (d) are known.

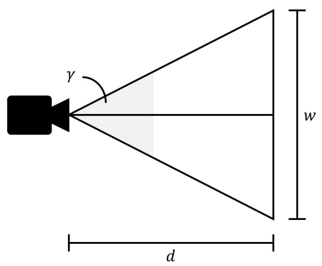

Figure 2.20: Schematic top view of camera indicating horizontal field of view (γ), depth (d) and width (w).

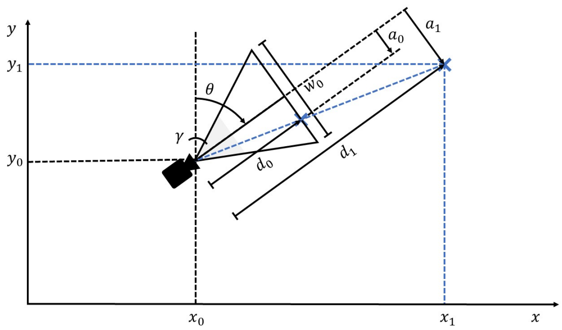

Figure 2.21: Schematic top view of an arbitrary point being projected on the view of the camera.

$$\tan(\gamma/2) = \frac{w}{2d} \quad \rightarrow \quad w = 2d\tan(\gamma/2) \tag{2.4}$$

This depth of the camera is found by setting up a system of three equations to uncover the x and y position of the camera (x_0, y_0) as well as the horizontal viewing direction (ϕ). These equations are set-up by investigating where a point at an arbitrary position (x_1, y_1) will be projected on the camera. The relevant dimensions for this problem are shown in Figure 2.21.

In this figure, a_1 and d_1 can be found by applying a 2D rotation to the relative x and y of the point to the camera by the viewing direction ϕ. This is seen in Equation (2.5).

$$\begin{pmatrix} d_1 \\ a_1 \end{pmatrix} = \begin{pmatrix} \sin(\theta) & \cos(\theta) \\ \cos(\theta) & -\sin(\theta) \end{pmatrix} \begin{pmatrix} x_1 - x_0 \\ y_1 - y_0 \end{pmatrix} \tag{2.5}$$

Due to rules of congruent triangles, the ratio of a_0 and d_0 must be equal to that of a_1 and d_1. We can combine this and rewrite to get an expression for the relative horizontal projection (a_0/w_0):

$$\frac{a_0}{d_0} = \frac{a_1}{d_1} \quad \rightarrow \quad \frac{a_0}{w_0} = \frac{a_1 d_0}{d_1 w_0} \tag{2.6}$$

Substituting in $d_0 = w_0/(2 * \tan(\gamma/2))$ and the expressions for a_1 and d_1 found in Equation (2.5), leads to:

$$\frac{a_0}{w_0} = \frac{a_1}{d_1} \cdot \frac{1}{2\tan(\gamma/2)} = \frac{\cos(\theta)(x_1 - x_0) - \sin(\theta)(y_1 - y_0)}{\sin(\theta)(x_1 - x_0) + \cos(\theta)(y_1 - y_0)} \cdot \frac{1}{\tan(\gamma/2)} \tag{2.7}$$

In this equation there are three unknowns (x_0, y_0 and θ). This means that three equations are needed. These three equations are retrieved by filling in Equation (2.7) with three points where both the relative horizontal projection (a_0/w_0) is known as well as the real world positions (x_1, y_1). For these positions, the bottom left sonic anemometer, top right sonic anemometer and the backwards sonic anemometer are used. It should be noted that for all of this it has been assumed that the field of view is known. Changing the field of view will change the position and viewing direction of the camera. The field of view is tuned to correspond with a ground truth distance and is found to be 40.75°. A visual representation of a camera fit on an image for the first of May is seen in Figure 2.22.

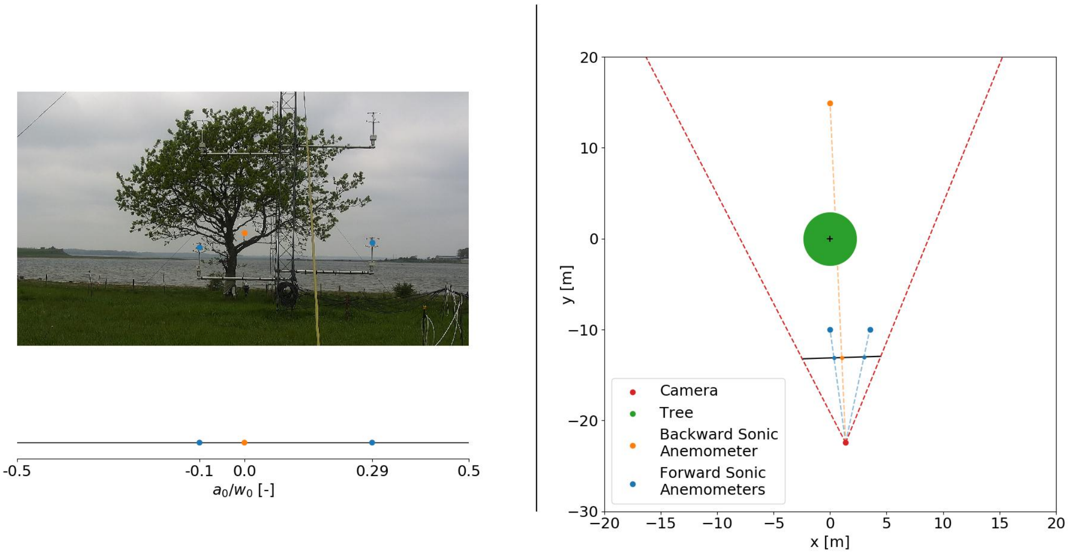

Figure 2.22: Visualisation of the relative horizontal projection (left) and the fitting of the camera to uphold these projections (right).

The distance from the tree found then can be used together with the field of view to discover the width of the screen at the depth of the tree by using Figure 2.20. The width found is divided by 1920, which is the number of pixels in the width to find the length of a single pixel. Squaring this length gives the surface area of a pixel, which can be used to transform the area in terms of number of pixels to meters squared.

2.2.6 Center of area

The tree is expected to sway with the wind. This swaying can be captured by the camera if the movement is tangent to the field of view of the camera. For this investigation, the horizontal swaying of the tree is of more interest, as it is expected to be more strongly related to the wind. This means that the movement of the tree in roughly the directions perpendicular to 290° can be noted.

Retrieving the center of area is done in three steps. First, the pixel coordinate of the center of all tree pixels is found. Second, the pixel coordinate of the base of the tree is found. Finally, the distance between these two coordinates can be converted from distance in pixels to distance in meters by multiplying by the pixel length. This allows for the horizontal center of area relative to the base of the tree.

The first step regards the pixel coordinate of the center of all tree pixels. This pixel coordinate is calculated per dimension by multiplying the dimension coordinate with the state of the pixel (either 1 if active or 0 if inactive) in the image I of size (Y,X). and then dividing by the total number of active pixels. This is seen for the horizontal dimension in Equation (2.8). A similar process is done for the vertical dimension.

$$ca_{h_{px}} = \frac{\sum_{j=0}^{Y-1} \sum_{i=0}^{X-1} I(j,i) \cdot i}{||\boldsymbol{I}||^2} \tag{2.8}$$

The second step requires finding the pixel coordinate of the tree. In this case the process will only be performed further than the first step for the horizontal dimension. The vertical dimension is expected to sway less and will thus not be investigated. This means that only the horizontal location of the tree needs to be known. By the camera fitting process described in Section 2.2.5 and knowing that the tree is located at (0,0) within that coordinate system, the horizontal pixel location of the tree can be extracted. The camera orienting process is done on a daily basis. Also important for this however, is the shift discovered in Section 2.1.3. This shift is recorded on an hourly basis and will thus require to be updated as such.

The final step is to look at the difference in the pixel location and multiply this by the pixel length which is updated daily. Furthermore, the change in pixel location of the base of the tree due to a camera shift is corrected for every hour. Expecting a drift of maximum 2 pixels, this means that the values recorded will not be influenced by camera movement by more than roughly 2 cm.

2.2.7 Optical porosity

The optical porosity (ϕ) is defined as the fraction of the enclosed area that is considered as void. In this case, this is considered as all pixels in the enclosed area that are not part of the tree. The optical porosity is thus calculated as shown in Equation (2.9) where S_t is the tree surface area and S_e is the enclosed surface area.

$$\phi = 1 - \frac{S_t}{S_e} \tag{2.9}$$

Again here it is noted that the total surface area is determined arbitrarily, as a length scale has to be chosen to dictate the circumference area, which directly influences the total surface area and thus the optical porosity.

2.3 Filtering

Some sources for erroneous or unusable data are present. In order to not include these observations, sources of the errors are investigated and filters are set-up which will filter out unusable observations. The process of filtering out data has seven layers. Only when an image makes it past a previous layer, will it be investigated for the next one. The seven layers are removing images that where not uploaded correctly, removing images taken throughout night mode, removing images where a mayor camera shift happened, removing images with relevant missing data, removing images with sun interference, removing images with lack of proper gradient and finally removing images where binarization failed.

Of these layers, the first layer filters files where the image has not been stored at all due to some error in connection. The other layers will be explained below.

2.3.1 Night mode

The camera has been set-up with purpose of monitoring the tree state more than two years ago. As a measure of doing so at night, a night mode had been added to the camera which takes pictures when the luminosity is no longer high enough to support normal camera functioning. An image taken in night mode can be seen in the left side of Figure 2.23.

Figure 2.23: A night mode image (left) and zoom in of the sky of that night mode image (right).

For the purpose of thresholding though, these images do not lend themselves well for being separated from their background, as there is only a single channel and hues in the tree correspond to those in the background. Furthermore, taking a closer look at the grayness of the sky (the right side of Figure 2.23), it is seen that rather than being a continuous gray color, it exists of specks of high and low value colors. This could be turned into gray by applying some form of Gaussian smoothing, but this would decrease the accuracy of the area measurement. Even if this was not the case, using a different channel combination for thresholding induces a slight bias, reducing the ability to relatively compare images. For these reasons any image which is seen to be grayscale (which only happens during night mode) will be excluded from the analysis.

2.3.2 Mayor camera shift / misposition

Periodically, the camera is moved from it's position or is rotated, this could lead to the inability to successfully apply all three steps of the masking process. The main limiting factor in the masking process is the fitting of the meteorological mast mask. Small shifts of up to 10 pixels are corrected. Beyond a shift, the camera is also seen to sometimes be mispositioned due to natural influences. Days where a shift of more than 10 pixels occurs or where the camera is mispositioned are manually filtered by inspecting the mean images of each day. An example of a camera shift is seen in the left side of Figure 2.24. An example of mispositioning of the camera is seen in the right side of Figure 2.24.

Figure 2.24: Example of mean images with a camera intraday shift higher than 10 pixels (left) or a camera that is mispositioned (right).

2.3.3 Missing data

Throughout the measurements, some images are incorrectly uploaded, leading to partially missing data. As pixels are uploaded from top to bottom, this missing data will manifest itself as lacking information at the bottom part of an image. An image with low amount and high amount of missing data is seen in Figure 2.25. Missing data is substituted with a single color (gray in this case). It is noted that an image can still be used for analysis, if only data is missing for parts of the image which are masked out. Thus, only images where missing data is present in the y-range where pixel-information is used are flagged as having missing data.

Figure 2.25: Representation of an acceptable amount (left) and an unacceptable amount (right) of missing data.

2.3.4 Sun interference

When the sun is in frame or close to being in frame, it will interfere with the detection of which pixels belong to the tree. An example of the sun interfering with the image is seen in Figure 2.26. In this figure, it can be clearly seen that at the upper left part of the tree and at the location of the streaks coming from the sun inhibit the ability of stating the class of these pixels (tree or air). The shows the need to filter out all images where the sun interferes with the binarization of the tree.

Figure 2.26: Example figure of the sun interfering with the image.

The performance measure for a method that will filter out these images is to minimize the false negatives while getting a recall of more than 98%. As a means of evaluating this, 256 images are binarized and labeled to check for sun interference. For half of the images, this is done at a time of day where the sun is expected to be in frame, and for the other half it is done where it isn't expected. In this data-set 70 images have sun interference and 186 do not. The reason that the split is not equally divided, is that clouds occlude the sun and thus no problem is found for processing some of the image taken at a time where the sun would otherwise interfere.

An initially proposed method is to note the fraction of fully saturated pixels (a maximum value in all three color channels, corresponding to white). When the fraction of fully saturated pixels are above a threshold, the image will be filtered out. This threshold is increased until the recall would decrease under 98%.

Applying this method to the data-set however, creates poor results, which is attributed to two reasons. Firstly, fully saturated rim-lighting is created around the tree due to a light-source behind the object. Secondly, fully saturated pixels can be present in clouds rolling over. This leads to the inability to make a good split as images without sun interference will also have fully saturated pixels. These two sources can be seen in Figure 2.27.

Figure 2.27: Example of an image (top) and their fully saturated pixels indicated in white (bottom) due to rim-lighting (left) and high value clouds (right).

To mediate this, three steps can be taken. Spatial averaging, temporal averaging and checking neighbouring images for sun-interference.

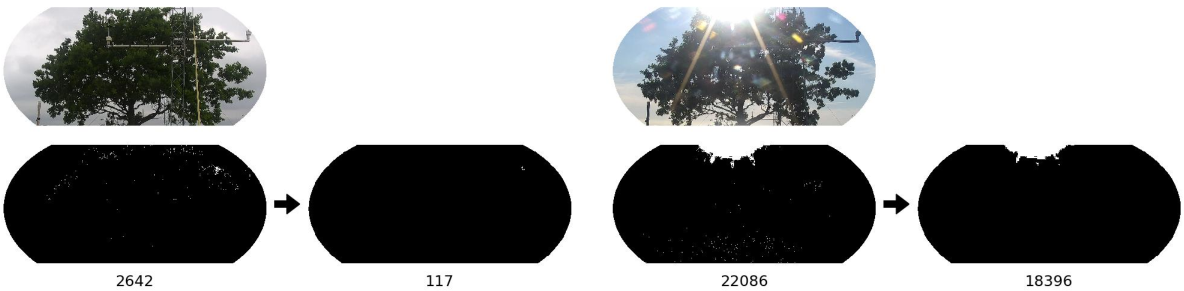

Figure 2.28: Two images, one with rim-lighting (top left) and another with sun interference (top right). Below the image the fully saturated pixels are indicated and to the right of that the fully saturated pixels after averaging in a 3×3 grid. Below each figure the amount of fully saturated pixels is shown.

Firstly, to mainly circumvent the problem of rim-lighting, spatial averaging can be applied to the image. This is achieved by down-sizing the image such that only if all pixels within an $r \times r$ region are fully saturated, it will be counted as fully saturated. The effect of this with a grid-size of 3×3 on an image with rim-lighting and an image with sun-interference can be seen in Figure 2.28. It is seen that more than 95% of saturated pixels for the rim-lighting image disappear, while the sun interference image still has a lot left.

Secondly, the sun is expected to be a somewhat continuous and stationary source of fully saturated pixels while clouds are not. Thus, to differentiate between clouds and the sun, temporal averaging can be applied. This is done by examining n consequent images in time and only considering a pixel to be fully saturated when the pixel remains so in all consecutive images. An example of this with a value of $n = 3$ for an image with fully saturated clouds and an image with sun interference can be seen in Figure 2.29. Again, the number of fully saturated pixels in the images that do not have sun interference is seen to be drastically reduced.

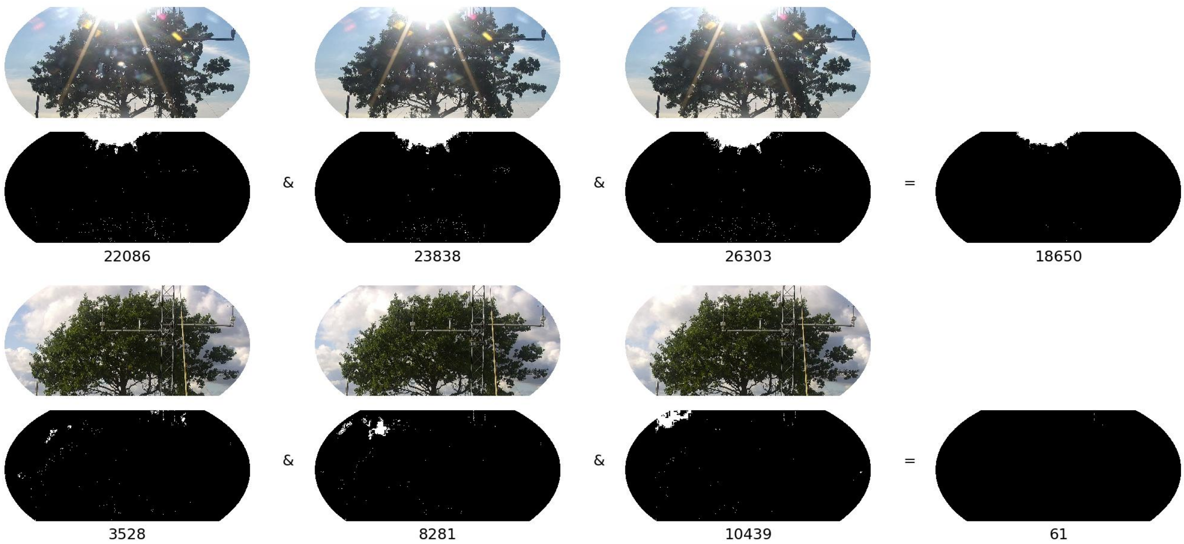

Figure 2.29: Two three-minute periods in time, one with sun interference (top) and one with clouds that have fully saturated spots (bottom). Below these images the fully saturated pixels are shown, and right of these images the result of only taking the continuous fully saturated pixels is shown. Below each binary figure, the amount of fully saturated pixels is shown.

Finally, due to occlusion of the sun, or entering/leaving of the sun from the screen, a case can occur in the temporal averaging where one image in the time-series will have no sun interference while the other ones will. For this reason, a neighbour check can be performed. This means that even if a time-series is seen to have continuous active pixels below the threshold, if their neighbouring time-series is above the threshold, they will still be counted as having sun interference to increase recall.

To check these three methods, five cases have been set-up with combinations of these methods. These five cases can be seen in Table 2.1.

Case	Spatial Averaging (r)	Temporal Averaging (n)	Neighbour check	parameters
1	X			r=[1,7]
2		X		n=[1,7]
3	X	X		r,n=[1,7]
4		X	X	n=[1,7]
5	X	X	X	r=3,n=[1,7]

Table 2.1: Description of the five cases investigated based on which methods they use and which parameter they investigate.

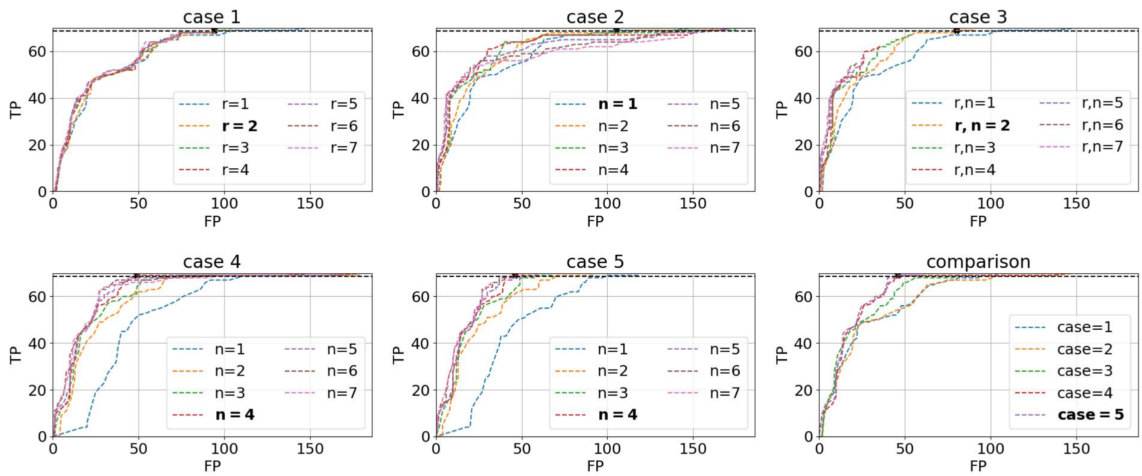

Figure 2.30: Display of the five aforementioned cases where True Positives (TP) are plotted against False Positives (FP). The best parameters per case are indicated in bold and the best parameters per case are compared in the bottom right figure.

Taking one of these cases as an example, putting the threshold at 100% of the pixels, will classify everything as negative (not experiencing sun-interference). This means that there will be no false positives but no true positives as well. Considering the case where the threshold is put at 0%, everything will be classified as positive (experiencing sun-interference). This means that the recall will be 100%, but the false positive rate will be too. Somewhere in between 0% and 100%, an optimum exists. Gradually increasing the threshold from 0 to 100, will increase both the true positive rate and the false negative rate. For each of these cases this process is performed, and the change in true and false positives is recorded, leading to the results seen in Figure 2.30.

From this figure it is seen that case 5 performs best when the recall is desired to be above 98%. It produces a false positive rate of 24.7%. Considering that 70% of this 24.7% is of images close to true positives in time, this result is accepted. It should be noted, that from Figure 2.30 it can be seen that the false positive rate can be decreased a lot by decreasing the required recall. For the purposes of this analysis however, a high recall is desired to minimize erroneous data-points in the final analysis.

Concluding, the method used to identify sun interference will include spatial averaging with a cell size of 3×3, temporal averaging over 4 minute intervals and will include a neighbour check. When more than 107 pixels (or equivalently more than 11 downsized pixels) are fully saturated in the final image or neighbouring image, all images in the 4-minute interval are flagged as having sun interference.

2.3.5 Rain interference / lacking gradient

Another source for erroneous data is rain drops on the lens. These can be visible as blurry areas that behave different from their surroundings or just distort the image as seen in Figure 2.31.

Figure 2.31: Example figure of the rain drops interfering with the image with blurry spots (left) and morphing of the tree (right).

Some way of filtering out images where these effects interfere with the extraction of the image aspects is required. The chosen approach to this problem is to train a logistic model to decide based on some image parameters, whether the image produces erroneous results due to rain interference. The reason for using a simple logistic model rather than using a neural network to do this task, is to gain some insights into how this classification will be made.

The image parameters that are given to this logistic regression are factors which are expected to deviate when rain interferes with the image. An artifact of the rain drop positioned on the lens, is a smoothing effect on the image, which is associated with a localized sharp decrease in the gradient of values in the image. This is seen when comparing the absolute gradient of an image with and without any moisture on the lens in Figure 2.32.

Figure 2.32: Visualisation the blurring effect of a rain drop on the lens by comparing an image with (right) and without (left) rain on the lens and inspecting the gradient of the image (bottom).

From this artifact of a lesser gradient, a new image is created which indicates a lacking gradient with respect to the mean image. This image is the Gaussian smoothed absolute gradient of the projected mean image minus the Gaussian smoothed absolute gradient of the projected observation. Why and how this is done is explained next.

Four steps are performed to achieve the lacking gradient indication image. First, it is noted that the tree will see more or less gradient based on how far away it is from the camera. To mend this, the area within the tree mask in projected onto a standard size of 500x800. Why this size was chosen is explained later.

Secondly, the absolute gradient of these images are taken to show where gradient is and is not missing. It is observed that the mean image inhibits a lower overall gradient than the observation. This is expected, as the mean image is the mean of the tree in many slightly different locations, which acts as a smoothing effect.

Next, this gradient is smoothed with a Gaussian filter with a σ of 10 pixels. This is done, as the final objective is a comparison between the observation and the mean image. The observation however, is expected to sway slightly in the wind, causing a slight mispositioning. To ensure good comparison and to get a heat-map of where gradients occur a lot in the image rather than a specific descriptor of the gradient image, this step is performed.

Finally, the observation is subtracted from the mean image, where negative numbers are set to zero. This means that there will only be a gradient left in the final image, at locations where there is a lower general gradient in the observation relative to the mean. The procedure is depicted in Figure 2.33.

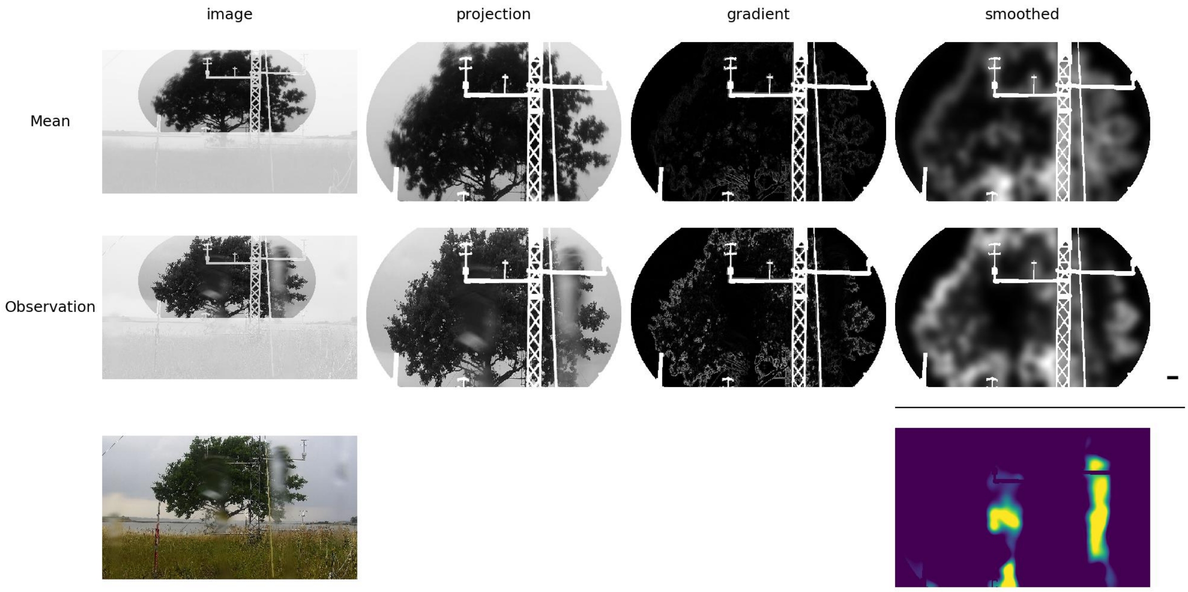

Figure 2.33: Visualisation of the steps taken to retrieve the lacking gradient indication image. On the top row, the mean image is seen, the row below is the observation. Bottom left shows the image and bottom right shows the lacking gradient image.

With this lacking gradient indicating image, some image parameters are identified which will lay the foundation for the learning process of the classification model. The attributes that are selected are: luminosity (mean value of the masked observation), mean and standard deviation of the gradient of the image and mean, standard deviation and maximum value of the lacking gradient indication image. All these attributes are recorded of 55 images where rain interferes with the lens, and 57 images close to these other images where rain does not interfere with the lens.

A model could be trained with these parameters, but three improvements can be made. Firstly, the model may be prone to over-fitting, thus regularization can be applied. Furthermore, the sigma used in smoothing and the projection size of the lacking gradient indication image can be optimized. Finally, not all features may contribute to improving the model, which is an extension of over-fitting.

To find the optimum of all these parameters and retain the ability to state the generalization error, two-layer k-fold cross-validation is used. For both the inner and the outer loop a value of $k = 5$ is used. The outer loop is split up in a test and training set, and the inner loop splits up this training set into a smaller training set and a validation set. The inner loop is used for finding the optimal model and the outer loop for evaluating this optimal model. An additional weight of $\boldsymbol{w}/C^2$ is added to the loss function as a means of regularizing. Here $\boldsymbol{w}$ is the norm of the weights and C is the inverse of the regularization strength. This means that for high values of C, the model will try to optimize while keeping its weights low. For C, 11 values logistically spaced between 10^{-5} and 10^{5} are used. As changing σ and the projection size will change the attributes of each observation, the options for these are limited to three options. The projection size will take on either (375x600), (500x800) or (625x1000) as a value and σ will take on a value of 5, 10 or 20, leading to 9 combinations of σ and projection size. Note that an aspect ratio of 1.6 is chosen for the projection size, to correspond to the aspect ratio of the ellipse in the masking process. It is also noted that all attributes are standardized by subtracting their mean and dividing by their standard deviation. This is to ensure that difference in scale does not create a bias in the training of the model.

For each of these 9 combinations, the regularization parameter is varied and the optimum is chosen. For one such combination, this can be seen in Figure 2.34. It is seen that by allowing the model to be more complex (increasing C), it is fit better to the training set. However, after a C value of 1, over-fitting occurs, and the validation set is predicted worse.

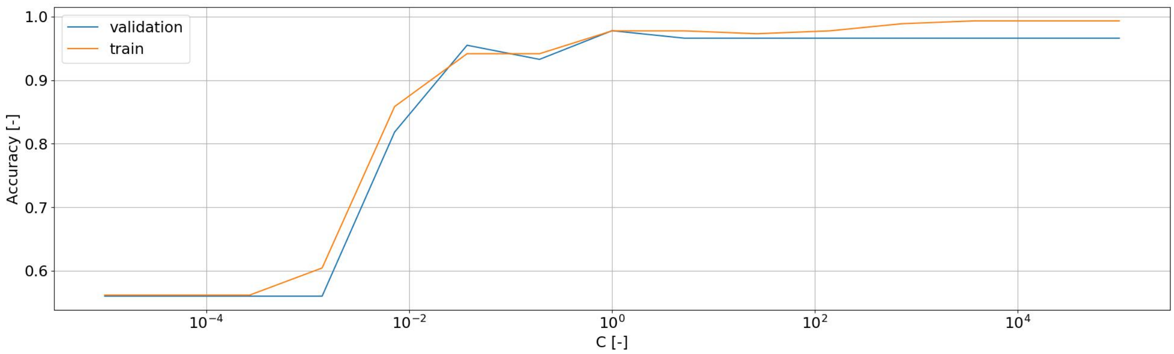

Figure 2.34: Example of regularization applied to the case of $\sigma = 10$, a projection size of (500x800) and using all features in the logistic model.

The optimum of all optimal validation performances found in each of the σ and projection size combinations then leads to the final optimal model. This entire process is used as a foundation for feature selection, and through backward and forward feature selection where the best performance or if equal, the lowest number of features is kept. The generalization accuracy of this model with the best feature selection is then once more trained over the outer training set and evaluated using the test set. The parameters chosen and the corresponding accuracy on the test set for all five outer folds is seen in Table 2.2. The features used for all five outer folds is seen in Table 2.3. It was found that of the optimal feature selections found in each fold, forward selection found better optima than backward selection.

Outer fold	σ	Projection size	C_{opt} [-]	Test accuracy [%]
1	10	(625x1000)	10^2	95.7
2	10	(500x800)	10^1	95.7
3	10	(500x800)	10^2	95.5
4	10	(500x800)	10^1	86.4
5	10	(500x800)	10^1	100

Table 2.2: Optimal model parameters found for each outer fold of the two-layer cross validation including the found accuracy on the test dataset.

Outer fold	Luminosity	Mean gradient	Std gradient	Mean lacking gradient	Std lacking gradient	Max lacking gradient
1				X		X
2						X
3						X
4			X			X
5						X

Table 2.3: Optimal features found for each outer fold of the two-layer cross validation. All features have been selected through forward feature selection.

Rather steadily it is seen that using a σ of 10 is considered an optimal solution, further it is seen that using (500x800) as a projection size is also found to be optimal four out of five times. As a complexity controlling parameter, C is seen to be best around 10 to 100. The accuracy on the test sets produce good results, averagely correctly estimating more than 94%. As for the features used, it is seen that many times, all but the maximum of the lacking gradient are discarded as being use full. In two cases it is accompanied by either the mean of the lacking gradient or the standard deviation of the gradient.

From this analysis it is concluded that most robustly, a model with a σ of 10, a projection size of (500x800) and a regularization parameter of 10 should be used. It is also found that exclusively looking at the maximum of the lacking gradient will produce good results, and lead to simplicity in the model. The model is trained using the entire data set and is used as a classification method for rain interference.

2.3.6 Failed binarization

The final filter that is applied is distinct from the other filters. This is as it is applied at the same time as the processing of the images rather than a pre-processing step. Failed binarization is noted as a split made in the histogram as part of the thresholding method which classifies a part of the sky as being part of the tree. This is usually the case when either a dark cloud rolls over which inhibits low color values, or when the sun sets and the sky will take on another hue. An example of a failed binarization due to both of these reasons is seen in Figure 2.35.

Figure 2.35: Example of failed binarization, caused by dark clouds (left) and a hue change in the sky (right).

Failed binarization is identified by looking at the outer rim (5%) of the masked area. In this area, nearly no pixels are expected to be identified as tree. If the fraction of pixels identified as tree is above a threshold of 0.05, the image is filtered for having unsuccessfully binarized the image. An indication of where this outer rim is located is shown in Figure 2.36.

Figure 2.36: Indication of the outer rim by the highlighted area, which indicates a failed binarization for a tree with dark clouds in the background (left) and a hue change in the sky (right).

CHAPTER 3
Results and discussion

The depiction of all results is done in three mayor steps. First, the data that has been filtered out to improve any further analysis is described. Second, the more general 6-month analysis is investigated. Finally, the 10-day specific analysis is investigated.

3.1 Filtering

Sequentially, seven steps are taken to filter. For the first step however, empty images are filtered out. This only pertains to less than 90 images out of more than 200.000. For this reason and for ease of depiction, these images are included in the second filter. The seven steps are sequentially described in Section 2.3 and will be referred to as "empty", "camera shift", "night mode", "missing", "sun interference", "lacking gradient" and "failed binarization". A time-wise indication of which images are filtered out at each step is seen in Figure 3.1.

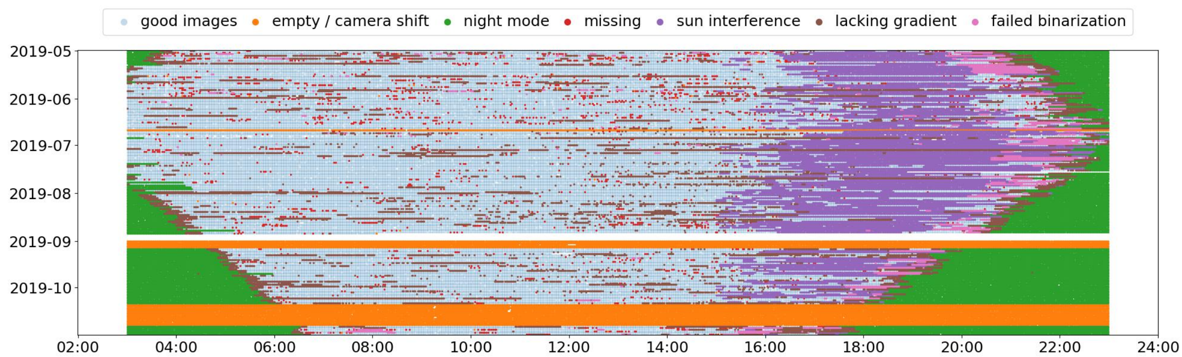

Figure 3.1: Images filtered in time where the x-axis shows the day and the y-axis shows the time in the day.

To investigate this figure sequentially, the first thing looked at is the first filter where a too big of a shift or mispositioning of the camera has occurred. For these it can be seen that specific days have been excluded from the research. The reason that many of these are sequential, is that something had occurred to the camera, which remained unnoticed for a while. As soon as the malorientation was noticed, images started being taken correctly again.

The second filter to be applied is the night time filter. The days can be seen to be truncated at the beginning and end based on whether the camera's night mode had been activated. The seasonality can be clearly seen where up until the 21^{st} of June (the longest day in the year), the days get longer, and afterwards, the days get shorter. The fluctuation in ending and starting night mode is due to night mode being luminosity controlled and not time-controlled. So, any heavily clouded days, will see the camera be in night mode more.

Thirdly, the images with missing data are removed. The cause for missing data is unidentified, but it is seen that some form of clustering beyond a random distribution occurs. This would indicate that there is some outside factor which influences whether an image would have missing data. A guess is that it could be internet connection related, and that at moments of high data-usage, problems occurred in the uploading of these images.

The fourth filter, filters out sun-interference. As expected, this only occurs later in the day when the sun sets, as the camera is roughly pointed towards the west. If all clouds would be absent, this should show a clear lower curve and upper curve, as the sun's temporal position in the sky will slowly drift. Any sporadic missing of this curve should be due to cloudy days. At some days it can be seen that there is no sun-interference at all. Inspecting images on these days confirms that this is due to a heavily clouded sky.

The fifth filter pertains to a lacking gradient. Two frequent occurrences of lacking gradients are seen. The first is just after the camera exits night mode and just before it enters it. This is as the night mode enters somewhat late, and before those moments, most contrast is already disappearing, inhabiting the proper division between sky and tree. The second occurrence that is seen, are small consecutive streaks on some days. Investigating these streaks show that the filter is doing one of its main tasks here; filtering out periods of rain.

Finally, the last filter, filters out failed binarization. This is due to the low-point in the histogram which is used for thresholding, being in the middle of the gradient of the sky. This occurs as the channel combination of 0.7 times the red channel, -1.24 times the green channel and 1 times the blue channel, depends on blue pixels being present in the sky. At sun-set, the sky will change hue, and will no longer get direct enough light to show blue. For this reason, it is seen that during sun-set, the final filter, detects this bad-split and filters out these images.

A quantitative overview of how many images are filtered out is shown in Figure 3.2. This is depicted in such a way that the percentage of filtered images relative to the images left from the previous step is indicated. From this it is seen that the sun-interference filter filters out a lot of images. It is known that many of these images are actually usable however, meaning that a better method of filtering out these images would be desirable.

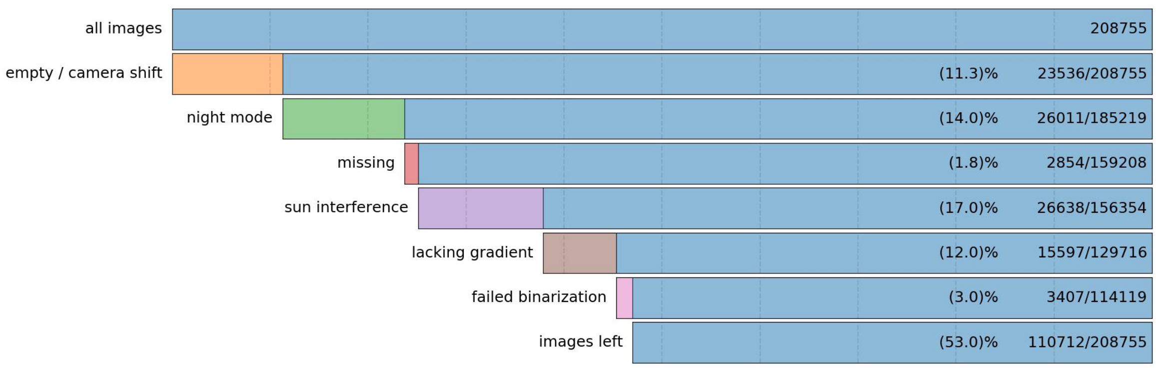

Figure 3.2: Visualisation of the number of filtered image per filtering step.

3.2 6-Month analysis

With an accurate map of which images can not be used for analysis, a period of six months is investigated. Analysing a longer data period gives more data, which means that more can be said about infrequent wind direction and speed ranges. However, using such a long period, will also include some seasonal dependencies such as leaf growth. Although these seasonal dependencies are interesting to investigate in and off themselves, they may also increase the difficulty of isolating relationships. For this purpose a lot of these statistics will be detrended later on in this chapter.

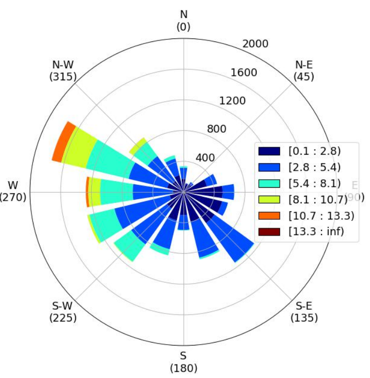

Figure 3.3: Wind rose showing how many 10 minute averages are available after filtering binned in wind direction and wind speed.

As an initial step however, all images within this 6-month period which are considered of good quality are analysed. The time series of all four extracted image aspects are shown to give an overview of the data. Next the relation between the wind speed and the tree surface area is investigated. Two trends are expected in the surface area. One of them is time-related and one of them is wind-speed related. In an attempt to isolate the wind-speed effects, the time-related effect on the surface area is estimated and the area signal will be detrended to isolate the wind-speed related effects. All usable ten-minute averages used have a distribution of wind directions and wind speeds as shown in Figure 3.3.

Comparing this to the wind rose of all unfiltered data in Figure 1.14, it is seen that southern directions are filtered out more heavily than other directions. This might imply that the filtering procedure had a bias towards filtering out southern wind directions. However, it is attributed to the dependency of the wind direction with the time of day. Mostly night-time and late-evening data is filtered out. The percentage of southern wind speeds (between 90 and 270 degrees) based on time of day, is seen in Figure 3.4, which confirms the previous statement.

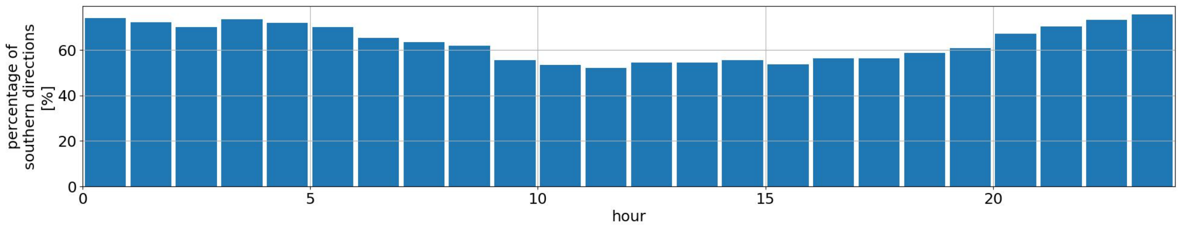

Figure 3.4: Percentage of wind coming from southern directions (between 90 and 270 degrees) at specific times of day.

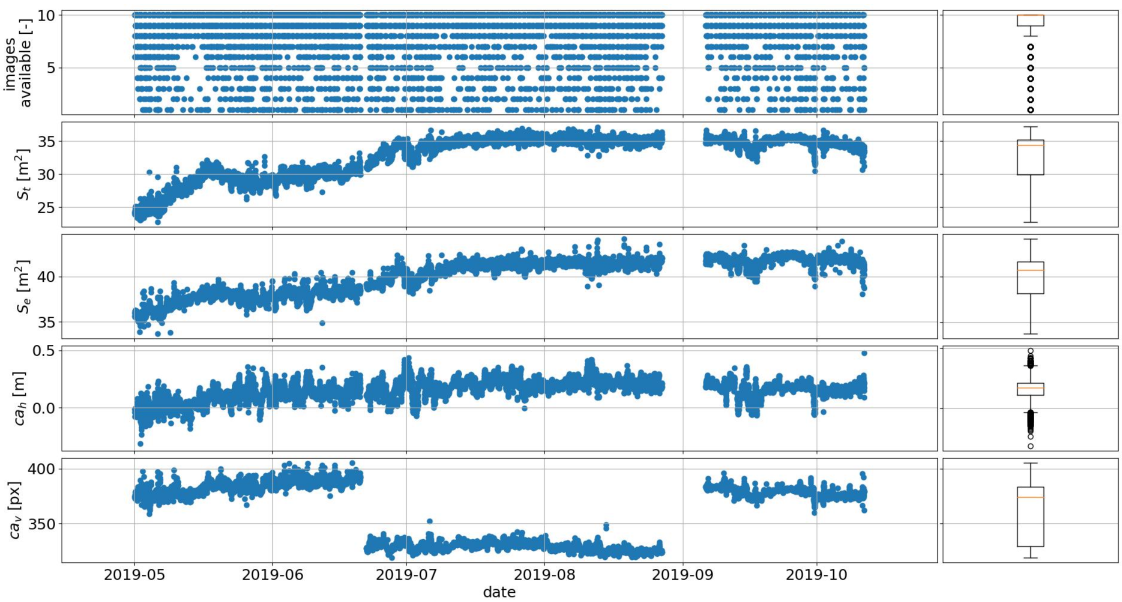

Figure 3.5: Overview of image aspects retrieved from good quality images, combined in 10-minute averages. At the right side a box-plot of the time-series can be seen. From top to bottom, number of images, tree surface area, enclosed surface area, horizontal center of area and vertical center of area is shown.

3.2.1 Overview image aspects

Suitable images are all collected in 10-minute statistics. Four image aspects are extracted. These are the tree surface area, the enclosed surface area, the horizontal center of area in meters and the vertical center of area in pixels. These and the number of available images per 10-minute average are inspected. All of these in the form of a time-series together with a box-plot indicating their distribution can be seen in Figure 3.5. Each of these statistics will be discussed.

Firstly, the number of images available for the 10-minute statistic. For validity of data, and for the 10-minute statistics to be accurate, at least a certain amount of images need to be present within 10 minutes. The distribution of 10-minute statistics based on the number of images within the interval is seen in Table 3.1. To retain more than 75% of the data, while upholding high availability, the threshold for usability of 10-minute statistics is set at having at least 9 images within it's interval.

number of images	10	9	8	7	6	5	4	3	2	1
percentage [%]	55.0	20.3	8.1	4.5	2.7	1.9	1.8	1.8	1.7	2.2

Table 3.1: Percentages of 10-minute statistics per number of images in the 10-minute window.

The tree surface area is seen to increase from May until roughly reaching a plateau around July. This is attributed to the growth of the tree. Furthermore, some unexpected strong fluctuations occur, which is attributed to the wind and will be discussed in Section 3.2.2.

Regarding previous years, the surface area was expected to increase around may and be fully grown by the beginning of June. After an increase in area from the beginning of May until roughly half-way through the month, a sudden decrease in surface area is seen. Which lasts for about a month, until it slowly increases and reaches its maximum value around the beginning of July. This deviating behaviour is linked to the bug infestation which has been described in Section 1.3 which was assumed to have started roughly half-way throughout May.

From Figure 3.5 it is also seen that the enclosed surface area also increases with time, though proportionally less than the tree surface area. This is as most of the seasonal growth of surface area is due to an increase in outline and mostly due to a decrease in porosity as leaves grow.

The horizontal center of area of the tree, is seen to slightly shift to the right by about 20 cm. It seems that the right side of the tree contains more branches and thus its surface area over there grows slightly more as leaves grow. Furthermore, heavy fluctuations in the range of 30-40 cm can be seen. This is expected to be due to heavy winds in perpendicular directions to the line of sight of the camera swaying the tree. This will be investigated in Section 3.2.3.

The vertical center of area is still depicted in pixels. The transformation of this statistic to actual height was deemed irrelevant as it's parameter is not expected to be linked to wind statistics. A dip can be seen in the middle of the time-series which is due to a shift in the camera at this time period.

3.2.2 Surface area and wind speed

Two dependencies are expected regarding the change in surface area of the tree over time. First, the tree surface area and enclosed surface area are expected to be dependent on the seasonality of the tree, as the tree will increase in size when it grows leaves in spring and will decrease in size as it sheds its leaves in autumn. Second, the tree is expected to decrease in size with an increase in wind speed, as it will start to streamline. These two dependencies are investigated here.

To begin, the tree and enclosed surface area are displayed below to the wind speed signal for a closer initial inspection. This is seen in Figure 3.6.

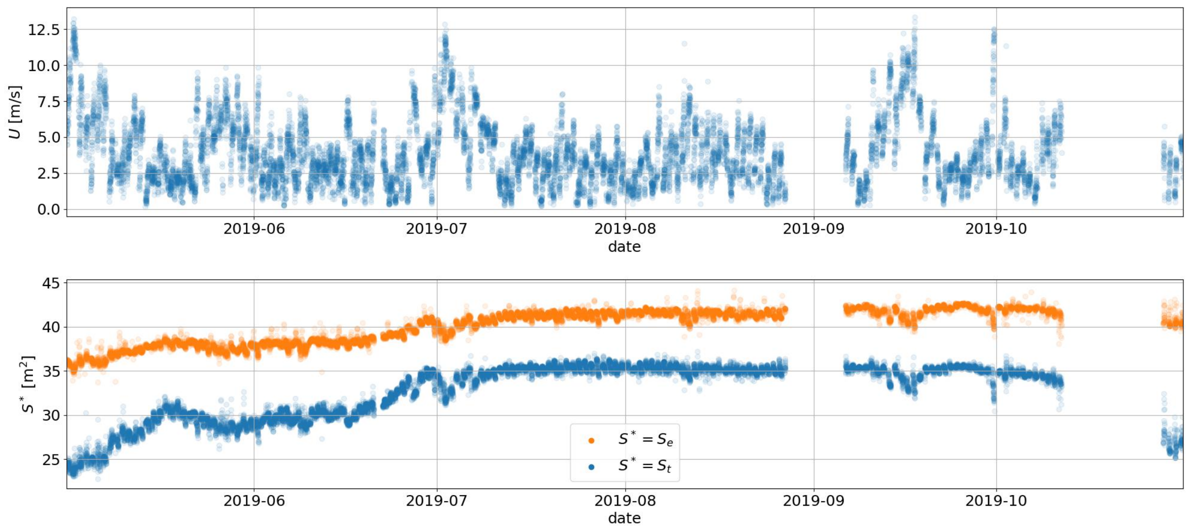

Figure 3.6: Time-series of the wind speed (top) and the tree surface area (bottom) where orange indicates the enclosed surface area and blue indicates the tree surface area.

From this figure both dependencies can already be seen. It is seen that the surface area increases from May to July, and that localized dips in the area graphs can be seen at points of continuous high wind speeds. It is noted that for all analysis however, the small period at the end of the time-series (at the end of October), is excluded. This is as there is not enough contextual data to create any trend for this data.

To better quantify how much the surface area decreases with a specific wind speed, the surface area can be plotted against the wind speed. However, due to the increase in area due to seasonality, this will not show a concrete trend as seen in Figure 3.7.

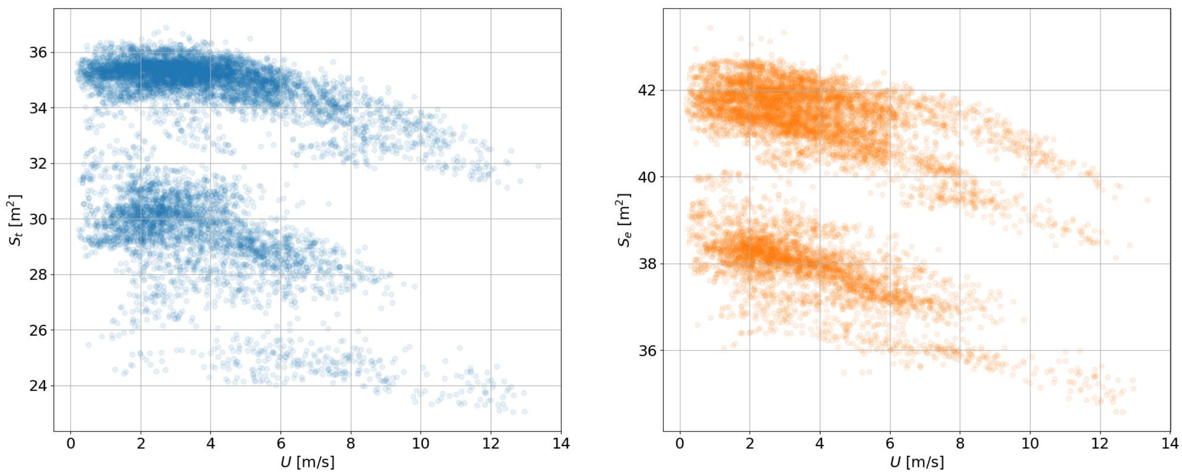

Figure 3.7: Tree surface area (left) and enclosed surface area (right) over wind speed for the period from the 1st of May 2019 to the 10th of October 2019.

Although no concrete trend is seen, some tendencies can be seen already. A point cloud which moves to the right, will also slightly move down, however, many of these point clouds exist, corresponding to different clumps of data when the tree saw growth. To create a somewhat more conclusive plot, rather than the surface area, the surface area over stationary surface area or also stated as the normalized surface area is inspected.

For this purpose the stationary surface area is described as the surface area without influence of the wind. This surface area is extracted from the data by applying a simple moving average to the data-points where wind speeds are lower than 2 m/s. This is applicable with the assumption that for wind speeds in this range, the surface area will only change negligibly relative to the 0 m/s condition. Continuing on this assumption, no wind directions are filtered out for the creation of this trend, as regardless of the direction, the wind is assumed to not influence the surface area of the tree in this regime. A value lower than 2 m/s could be chosen to make these assumptions more valid, but this would come at the cost of available data for creating a trend. As a balance between validity and availability 2 m/s was chosen.

This trend through the tree surface area and enclosed surface area, together with both areas plotted over wind speed is seen in Figure 3.8.

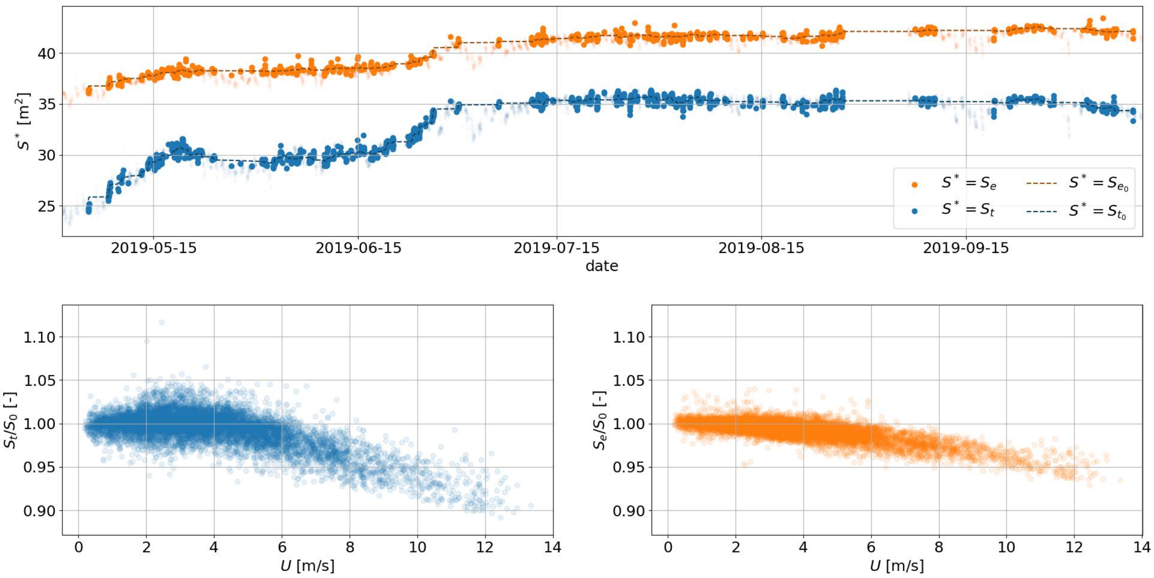

Figure 3.8: Detrended tree surface area over wind speed (bottom) for the tree surface area (left) and the enclosed surface area (right). Above this a visualisation of which points (boldly printed) are used to create a evolving stationary surface area (dashed lines) by means of a simple moving average with a window size of 40.

Figure 3.9: Examples of the tree in low leaf condition (left) and a high leaf condition (right) where the tree is seen to be less porous and bigger in the high leaf case.

Now in the detrended data, the decrease in surface area with increase in wind speed becomes more clear. It is seen that the tree surface area decreases more than the enclosed surface area, which is expected, as an increase in optical porosity is expected with an increase in wind speed. Furthermore, it is seen that a higher fluctuation in relative tree surface area is seen than that of the relative enclosed surface area. This is attributed to the overestimation of gaps in a brightly illuminated tree. When the tree is brightly illuminated, it will see bright spots on the leaves which the thresholding method will identify as sky. This means that the tree surface area has an added dependency on the illumination of the tree. This problem does not translate to the enclosed area of the tree, as it does not influence the outline which is its determining factor.

In these correlations however, no distinction is made between the rigidity of the tree (high vs low leaf case) and between the wind direction. The influence of these two factors are investigated. A high leaf and a low leaf case are displayed in Figure 3.9 to exemplify these two conditions. It can be seen that the high leaf case is less porous and in general has a bigger crown.

For a more rigid tree, the decrease in surface area with wind speed is expected to be less. This would mean that a tree with less leaves would experience this decrease in surface area less than a tree with more leaves. To confirm this, the points in the relation clouds in the bottom of Figure 3.8, belonging to a month with less leaves and a month with more leaves are investigated separately. The 20$^{\text{th}}$ of May until the 20$^{\text{th}}$ of June is taken as the month with a low amount of leaves and July is taken as the month with a high amount of leaves. These months are chosen as the stationary surface area in the first month is seen to not yet have plateaued yet, while the one in July has. These months are also chosen for their high availability, somewhat stable stationary surface area and large range of wind speeds. The result of this analysis is seen in Figure 3.10.

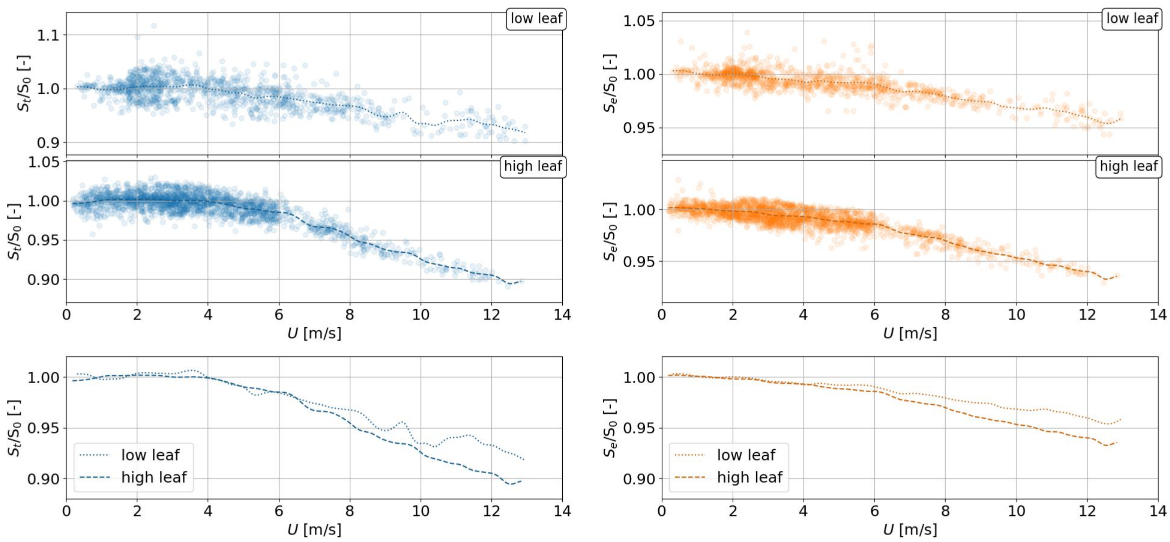

Figure 3.10: Separate inspection of tree surface area (left) and enclosed surface area (right) of a low leaf month (top) and a high leaf month (middle). Moving averages are compared in the bottom of the figure.

It is noted that the low leaf case, still had leaves, though less than the high leaf case. However, this figure does show that leaves account for a large part of the streamlining of the tree.

Besides investigating the relative surface area change based on leaf condition, a wind direction based analysis is performed. Data points are binned in nine 40° wind direction bins centered around 20° with an increasing offset of 40°. The result of the moving average of all data within these bins is seen in Figure 3.11. It is noted that for this figure, that some speed range and wind direction combinations contain sparse amounts of data.

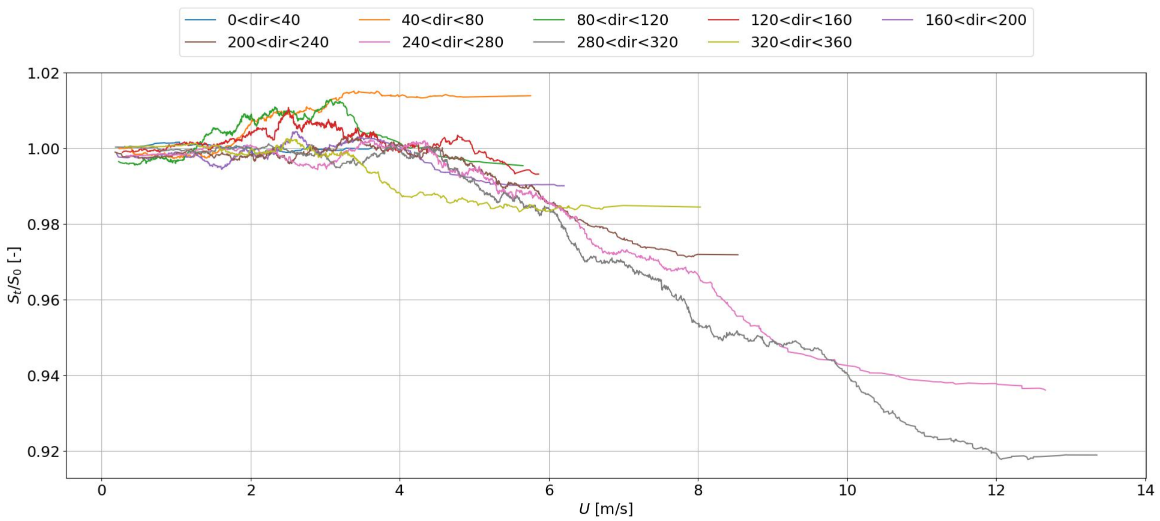

Figure 3.11: Simple moving average with window size 40 of the detrended tree surface area over wind speed for 9 different wind direction bins.

It is seen that based on wind direction, the graph changes significantly. While for wind from 160 to 360 degrees, the surface area decreases with an increase in wind speed, for wind directions between 40 and 160 degrees, an increase in surface area is noted for the low wind speed regime.

To explain this, a more broad view has to be taken as to why the tree streamlines in the wind. It is beneficial for a tree to decrease loads in high winds, to reduce fatigue or general damage. When looking at the wind-rose in Figure 3.3, it is seen that most high wind speeds come from the west. It would thus be most beneficial for the tree to grow a directional advantage in streamlining in this direction. The camera is roughly centered around 290 degrees, meaning that only wind directions near this direction and diametrically opposed to this direction depict the frontal surface area. This means that the frontal surface area is well represented by the wind directions centered around 300 and 100 degrees. As has been stated, it is seen from Figure 3.11 that wind from the 300 degree centered sector slowly decrease the surface area with wind speed, and winds from the 100 degree centered sector first increase surface area before decreasing it. An analogy with an umbrella is made. If a strong wind would come towards the top of the umbrella, its surface area would shrink, if the wind comes from the stem of the umbrella however, due to the range of motion of the umbrella, it would first grow before shrinking once more.

This difference in trend per wind direction is attributed to the more frequently and intensely occurring western wind speeds which would lead to the tree growing to be more able to streamline in this direction. Apparently, efficient streamlining in one direction may come at a decreased efficiency and possibly even an adverse effect of streamlining in another.

3.2.3 Wind direction and horizontal center of area

Relations between the horizontal center of area and the wind direction are investigated. The two signals are shown in the left side of Figure 3.12. In the right side they are plotted against one another.

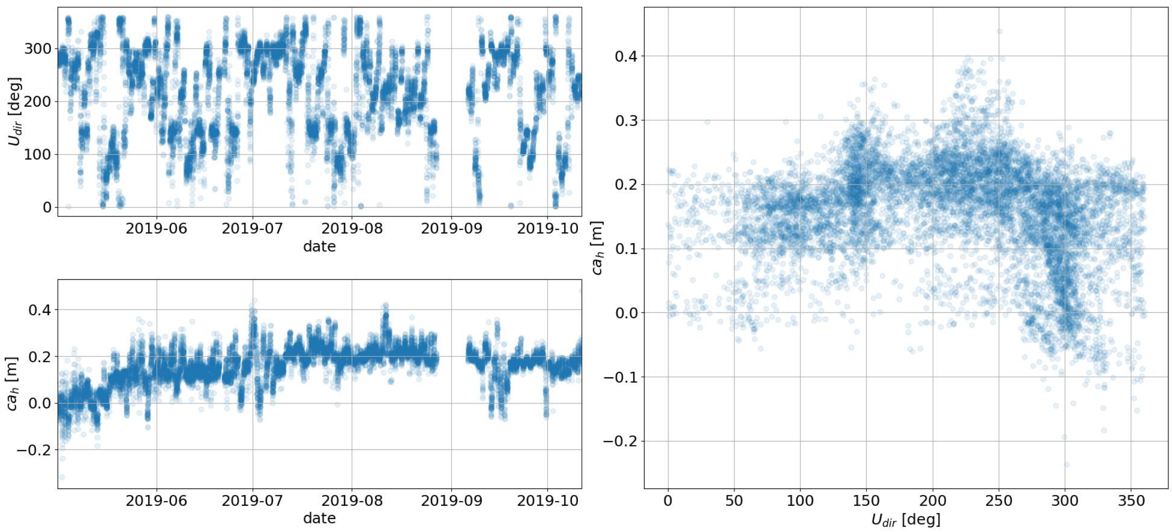

Figure 3.12: Time-series of the wind direction (top left) and the horizontal center of gravity of the tree (bottom left) and their relationship (right).

Looking at the bottom left part of Figure 3.12 which shows the horizontal center of area time series, a general trend can be seen. It slowly increases as the tree grows. A tree without leaves is expected to have a different center of area than that of the same tree with leaves. Thus, as the tree grows leafs, it follows that the center of area will shift towards that of the full leaf condition.

Looking at the horizontal area versus wind direction however, hardly any relation can be seen. This is as the wind direction may dictate the direction and projection of swaying, but the wind speed will determine the magnitude. Secondly, the previously mentioned drift in center of area also makes it harder to find any trend in the data.

To look at the shift in horizontal center of area rather than the horizontal center of area itself, the horizontal center of area relative to a moving average of itself at wind speeds below 2 m/s is taken. For this to be valid, the assumption is made that the horizontal center of area does not shift significantly when wind speeds are lower than 2 m/s. The normalized horizontal center of area (center of area c_h minus the stationary center of area c_{h_0}) is shown in Figure 3.13.

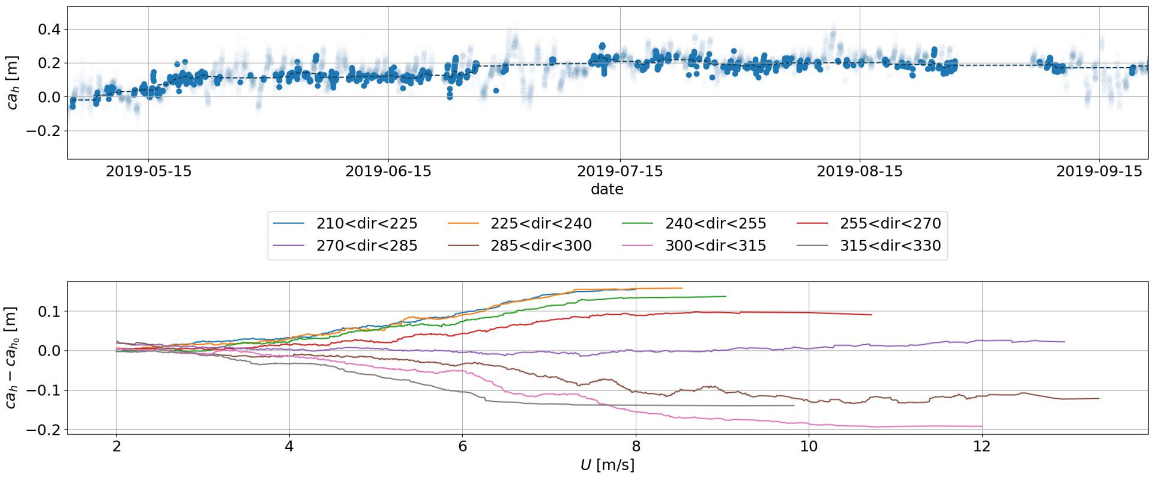

Figure 3.13: Time-series of the wind direction and the horizontal center of gravity of the tree detrended by applying a simple moving average with a window size of 40.

It is seen here that the deflection increases with wind speed. Assuming a simple cantilever beam relationship and a point force, the deflection is expected to be proportional to force, and while the slope of the trend seems to increase initially, it flattens out after a while. It is also seen that the change in deflection direction occurs roughly around 280 degrees. This is attributed to the non-homogeneity of the tree, as otherwise this change would be expected around 290 degrees.

Another interesting observation is the flattening out of the curve of directions between 315 and 330 at a low wind speed compared to curves at slightly lower angles. The reason for this flattening out behaviour could be due to the tree existing out of many objects with different stiffness. Low stiffness objects (leaves) allow for a swift response in the lower speed region, and medium stiffness objects (small branches) take over after that. Finally, in the higher speed regime, the small branches have almost all maximally deflected, and the big structure of the tree is so stiff that it only responds minimally to the wind.

3.2.4 Porosity

From Figure 3.8 it is seen that the enclosed surface area and the tree surface area do not change at the same rate. This is directly related to a non-constant porosity. The porosity can be plotted as a function of time. The stationary porosity is extracted by applying a simple moving average to all data points associated with wind speeds below 2 m/s. This then allows for the relative porosity to be compared with wind speed. Both the detrending process and the relative porosity versus wind speed is seen in Figure 3.14.

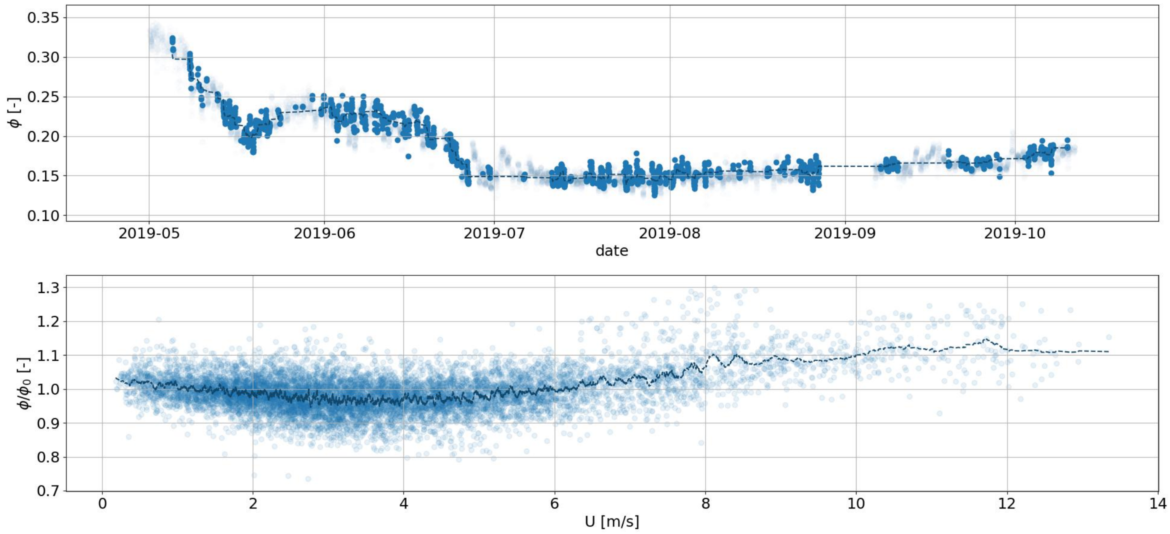

Figure 3.14: The porosity time signal (top) with a simple moving average with window size 40 as a dark dashed line through points associated with wind speeds lower than 2 m/s indicating the stationary porosity. Also, the relationship between the relative porosity and the wind speed (bottom) with a simple moving average with window size 40 shown as a dark dashed line.

It is seen that for low speeds, the porosity decreases with an increase in velocity. This is in line with the observations made in Figure 3.8 where the total surface area decreases more quickly and after roughly a wind speed of 4 m/s, the tree surface area decreases more quickly. This can be linked to the tree first presenting itself as a somewhat nonporous structure, making most deformations happen at the edges of its crown, reducing enclosed surface area, and after a higher wind speed, letting through more wind and becoming more porous.

3.3 10-Day analysis

With an overview of the 6-month period, a more specific period of 10 days is investigated. These 10 days range from the 2nd of July to the 12th of July 2019. During this 10-day period, the temporal resolution will be changed from 10 minutes to 1 minute. Furthermore, a directional filter has been applied. Only observations with wind directions between 280 and 300 degrees are considered. The goal of this analysis is to discover the influence of surface area on the proportionality between the force and the wind speed. Furthermore, a relationship between the optical porosity and the wind deficit is investigated. Finally, the importance of the filtering process is evaluated. Before these three analyses are performed however, any time shift in the recording of images is investigated.

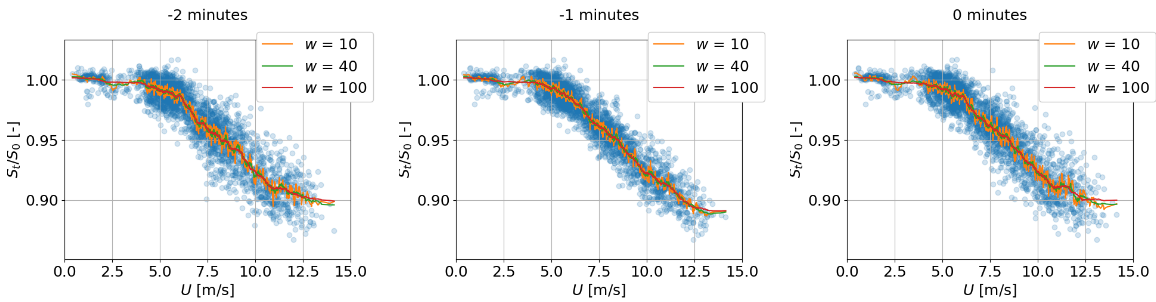

Figure 3.15: Three examples of the area vs. wind speed with a simple moving average for time shifting the wind data -2 minutes (left), -1 minute (middle) and not shifting the wind data (right).

3.3.1 Time shift

It is not certain that the time-stamp of the images is concurrent with the time at which the image is taken. This is investigated by assuming that the strongest relationship exists between the wind speed and the surface area of the tree when there is no time-offset between them. A time-offset, will thus lead to a weaker relationship. To investigate this time-offset, the surface area of the tree is compared to wind statistics of its timestamp with an offset ranging from minus ten minutes to plus ten minutes. A visual inspection, would show a decrease in spread for the introduced offset which matches the actual offset between the image and the timestamp. As an example, the area plotted against the wind speed for three different offsets is seen in Figure 3.15.

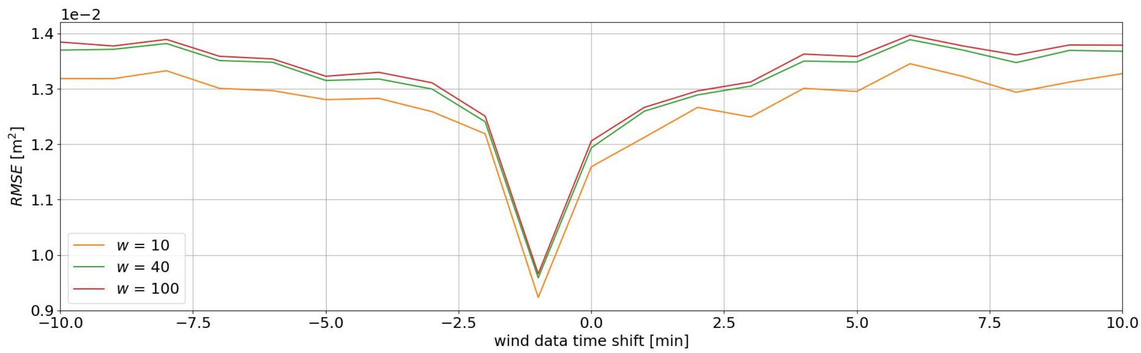

Figure 3.16: Graph of the RMSE w.r.t. three simple moving averages for time shifts between minus 10 and plus 10 minutes.

A way of quantifying the spread of this data, is to first calculate the simple moving average of the data-set. As the window size (w) will influence the fit, this process is done for a w of 10, 40 and 100. With these moving averages, the root mean squared error (RMSE) of the data-set relative to the fit is calculated. This can be seen for all three values of σ in Figure 3.16.

It can be seen that regardless of window size, a minimum spread is seen when a time-offset of minus one minute is applied to the wind statistics. This would mean that there is an offset of roughly one minute between the moment when the picture is taken and when it's timestamp is recorded given that the time stamps of the meteorological mast are true. The rough symmetry between the RMSE of one minute before and one minute after the optimal offset for all values, is a good indication that the offset is close to minus one minute. For future results, an offset of minus one minute will be applied to all non-image related statistics to mediate the dissonance between the timestamp and the image. For all previously presented results including those in the 6-month analysis, this offset has already been added.

3.3.2 Effect of surface area on F-U curve

The force on the tree has been estimated by measuring the bending moment as described in Section 1.3 and can thus be related to the wind speed. As has been described in the introduction, the drag force over wind speed curve has been described by Vogel, 1984 to be of the form $a \cdot U^{2+E}$ for a higher speed regime. The location of this higher speed regime is roughly around 6 m/s, but varies based on the tree and season. The value for E for this specific tree around the time of investigation can be seen by fitting a curve of the form $a \cdot U^b$ to the force over wind speed curve. This is shown in Figure 3.17.

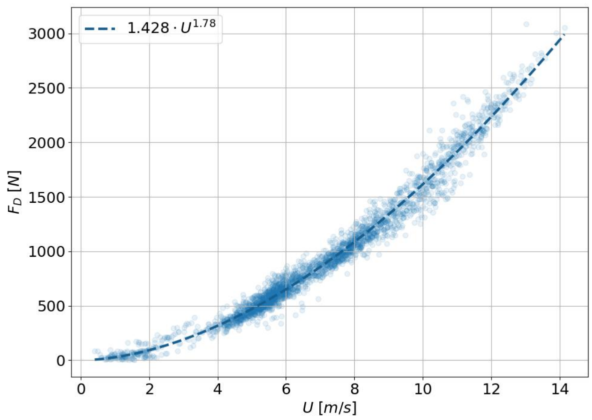

Figure 3.17: Drag force exerted on the tree over wind speed, with a power curve fitted on it showing the deviation of drag force from a U^2 dependency.

From this figure it would seem that indeed, the drag force does not increase proportional to the velocity squared. Rather, a value of E equal to -0.22 is found. The fact that this value is non-zero would either assume that the drag felt by the tree can not be described to be purely inertial in nature or that factors in the drag equation mentioned in Equation (1.5) (repeated here in Equation (3.1) for convenience) which have been assumed constant with wind speed are not.

$$F_D = \frac{1}{2}\rho U_\infty^2 S C_D \tag{3.1}$$

This would mean that either ρ, S or C_D would vary with wind speed. The density is not expected to vary much with wind speed. This is confirmed by plotting it against wind speed in Figure 3.18. Here only a weak relationship seen. Furthermore, if this effect was to be expressed in the form of a contribution to the Vogel exponent, this contribution would be positive rather than the observed negative Vogel exponent.

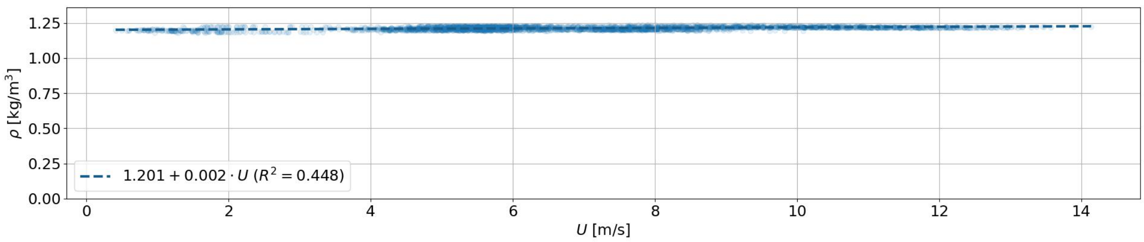

Figure 3.18: Density of the air over wind speed with a linear curve fitted to it. The density is seen to have little dependence on the wind speed.

This leaves a reducing surface area and drag coefficient with wind speed as leading suspects for a negative value of E. This coincides with what Gosselin, 2019 has stated, which is that three leading causes can be identified for the negative value of E. In addition to the two previously mentioned sources, it was also suspected that the porous nature would allow some of the wind to pass through the tree and make the tree experience a lower equivalent velocity. The effect of porosity is included by both considering the tree and enclosed surface area in further analysis.

Of all these three factors, the surface area is readily investigated with the data extracted from the images. As the stationary surface area of the tree is still slightly increasing in the beginning of the 10-day analysis, this would influence the relationship between surface area and wind speed. For the purposes of investigating the change in surface area with wind speed, the surface area is normalized with respect to the stationary surface area determined in Section 3.2. The normalized surface area over wind speed can be seen in Figure 3.19. Here it is seen that indeed, with an increase in wind speed, the surface area decreases, which would lead to a negative value of E.

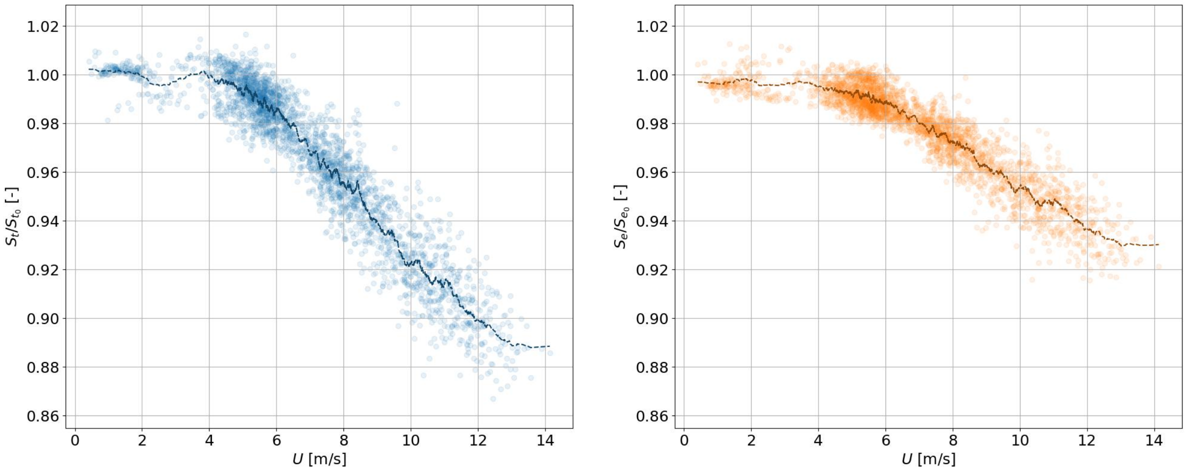

Figure 3.19: Normalized tree surface area (left) and enclosed surface area (right) over wind speed. A moving average (dark line) with window size 40 is applied to the data.

Interesting about this graph, is that no fit is found which describes the entire wind speed regime well. In the low speed regime (up until 6 m/s), the decrease in surface area seem to build up, until stagnating to something somewhat linear. For Vogel's expression to hold however, Equation (3.2) should also hold.

$$\frac{1}{2}\rho S C_D \propto U^{-E} \tag{3.2}$$

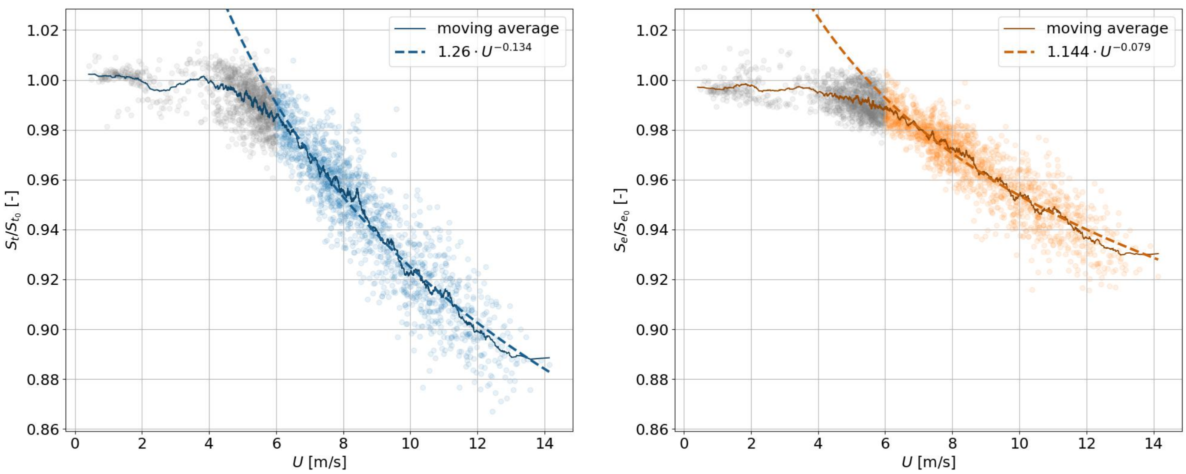

Figure 3.20: A curve of the form $a \cdot U^b$ fitted to the high speed regime (above 6m/s) of the tree surface area (left) and enclosed surface area (right) over wind speed.

This means that it is not unlikely that the surface area will closely follow the form $a \cdot U^b$ where b is negative. Such an expression would however tend towards infinity as the wind speed approaches zero, meaning that it will clearly only hold for the higher speed regime, as the tree can not have an infinite surface area. A curve in this form is fitted to all data with a wind speed higher than 6 m/s. The result of this is seen in Figure 3.20.

The fit is seen to follow the data well, though it overestimates the area above speeds of 11 m/s.

Another way of visualizing the E is by non-dimensionalizing the drag force with velocity squared and considering the Vogel exponent to manifest itself in the drag coefficient. In other words, assuming that C_D will hold all the proportionality of U^{-E}. Doing so by non-dimensionalizing the C_D with the stationary surface area leads to Figure 3.21. It is noted that as C_D is calculated by dividing by U^2, small fluctuations will cause dis-proportionally high values of C_D for really low wind speeds. For this reason, no values below 4 m/s are shown as these exhibit a large spread of C_D values which inhibits the proper investigation of the higher speed regime.

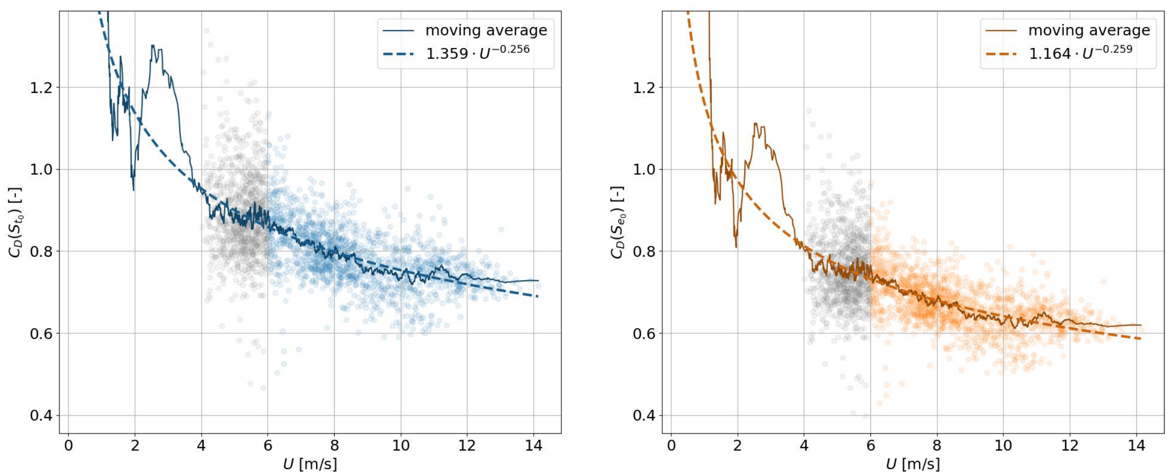

Figure 3.21: The stationary drag coefficient C_D as a function of wind speed with a curve of the form $a \cdot U^b$ fitted to the high speed regime (above 6m/s).

This shows a value for E of roughly -0.256 for the stationary tree surface area case or -0.259 for the stationary enclosed surface area case. This value is different from the value found in Figure 3.17 for three reasons. First, this value is only fitted to data points where wind speeds were higher than 6 m/s. Second, the change in stationary surface is now taken into account. Finally, although it has a smaller contribution, the density is incorporated in this graph as well, which was seen to have a minor relationship with wind speed.

The reason that both the stationary tree surface area and the stationary enclosed surface area showed different Vogel exponents is due to the slightly different trend in both evolving area's over time.

When non-dimensionalizing C_D curve with the dynamic surface area, the drag coefficients related to higher wind speeds are expected to increase. This is as drag coefficient is inversely proportional to the surface area, which will decrease. The C_D over U curve for dynamic surface area is seen in Figure 3.22.

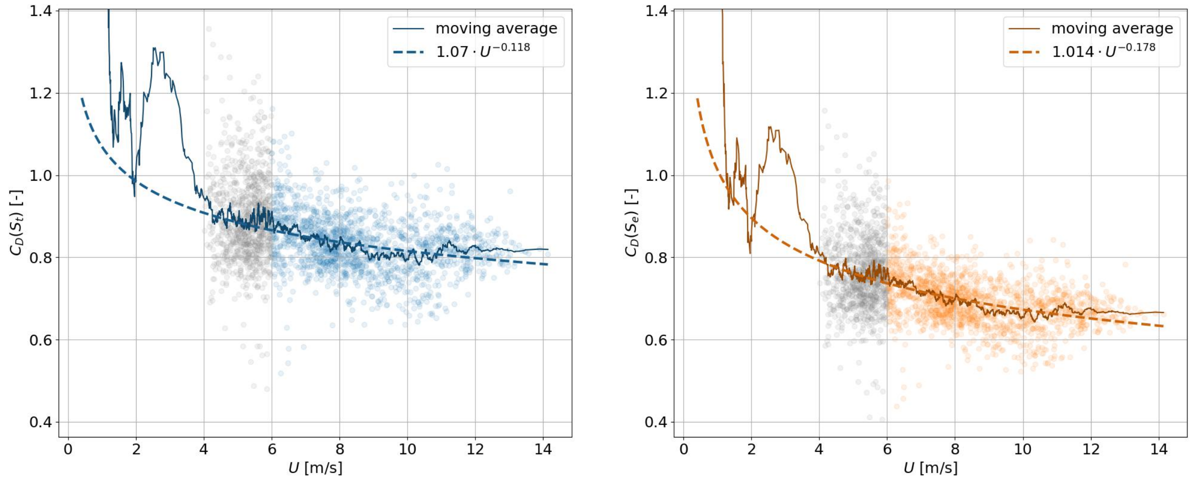

Figure 3.22: The dynamic drag coefficient C_D as a function of wind speed with a curve of the form $a \cdot U^b$ fitted to the high speed regime (above 6m/s).

The contribution of the surface area including the porosity can be seen from the left figure, as it is equal to the difference in exponent to the Vogel exponent found in Figure 3.21. This is a contribution of 0.138. This roughly corresponds to the value found in Figure 3.20. The reason that these numbers don't completely coincide is because a small relation exists between the noise of both the force measurements and the wind speed measurements. Excluding the effect of porosity, the exponent of the enclosed surface area is to be compared. For this, a contribution of 0.81 is found.

These results are compared to a Red Adler (Alnus Rubra) tree sapling investigated in Vollsinger et al., 2005 which is most analogous to the European Oak tree in this investigation. The Red Adler saw a decrease in C_D of roughly 50% for the stationary area case and 20% for the dynamic area case between a wind speed of 4 m/s to 14 m/s. In the case of this European Oak tree, a decrease in C_D of 25% and 15% were seen respectively for the stationary and dynamic surface area case within the same wind difference interval. This could be attributed to the difference in scale between the two studies or between the differences in the two trees.

3.3.3 Porosity and wind deficit

The wind deficit is linked to the porosity (ϕ). An indication of wind deficit is the downstream wind speed relative to the inflow wind speed or U_d/U_u. When this value is 1, no energy has been extracted from the wind. When this value is 0, all energy is extracted, and wind has halted. With an increase in porosity, this value is expected to increase, as the tree will let through more air. This value has been shown to be highly sensitive to the wind direction by E. Dellwik et al., 2019. A somewhat stable wind direction interval is to be chosen, where U_d/U_u is somewhat constant for all wind speeds. This stable region is found by plotting the wind deficit values over the direction for several wind speed bins. This is shown in Figure 3.23. Note that wind speeds are shown from 2 m/s onward, as below this range, the wind deficit value becomes highly variant.

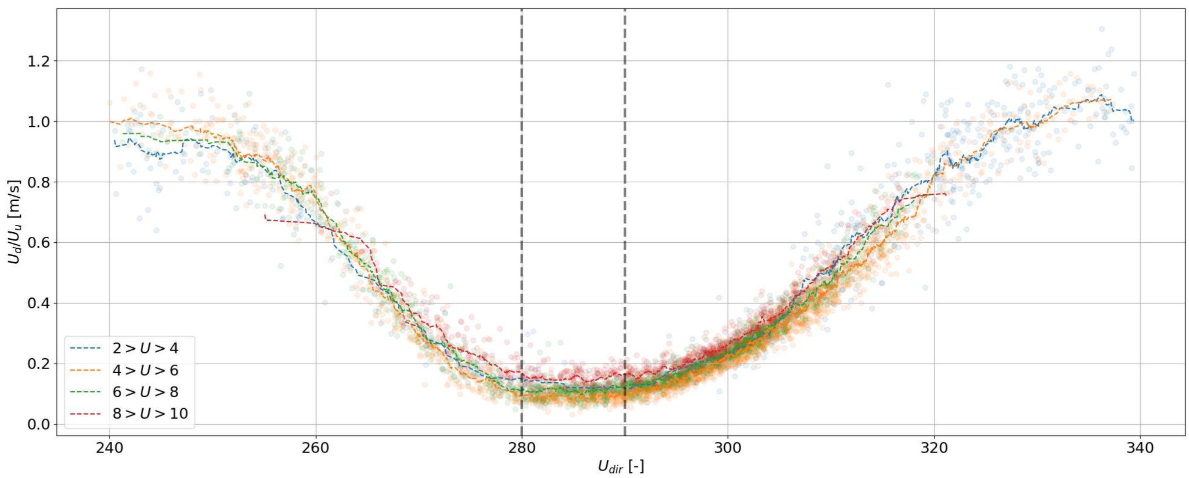

Figure 3.23: Wind deficit over wind direction for four different speed intervals. A section between 280 and 290 degrees is indicated with black dashed lines.

First of all, it is seen that this curve is not symmetric as the tree is non-homogeneous. Second, a similar trend is seen for most wind speed ranges, where higher wind speeds let through slightly more wind. From this, it is seen that roughly between 280 and 290 degrees, the wind deficit is stable.

Using a directional filter between 280 and 290 degrees, the wind deficit as a function of porosity is plotted. In this case the actual porosity is considered rather than the porosity relative to the stationary porosity, as in Figure 3.14 it is seen that the stationary porosity does not change much during the 10-day analysis period. The wind deficit as a function of porosity is seen in Figure 3.24.

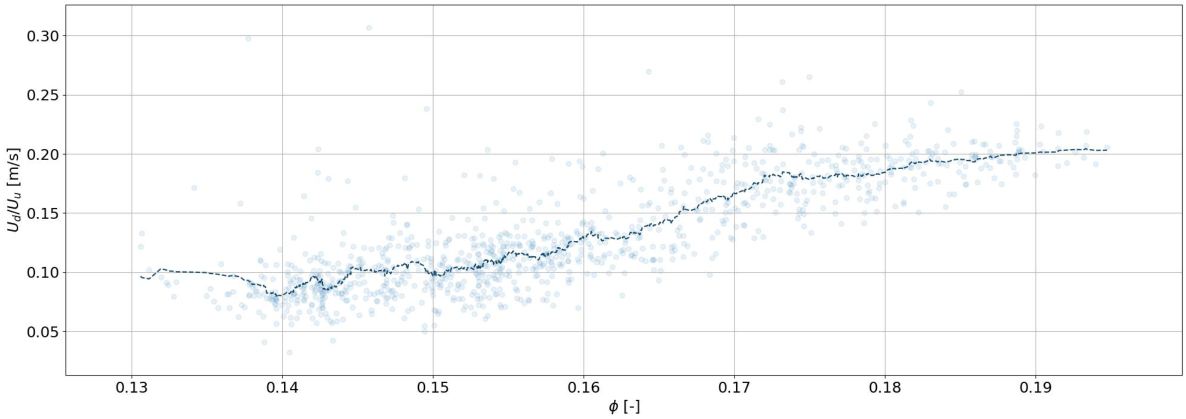

Figure 3.24: Wind deficit over optical porosity for wind directions between 280 and 290 degrees.

It is seen that indeed with an increase in porosity, the wind that is let through the tree increases. However, this relationship is not fully linear. This could be attributed to more skin friction drag being present, flattening off the curve more for higher porosity.

3.3.4 Filtering Evaluation

During the 10-day analysis, as a manner of evaluating the effectiveness of the filters, also filtered out images have been processed. This has been done for four filters, the missing data filter, the sun-interference filter, the lacking gradient filter and the failed binarization filter. The other filters were deemed to produce too much of an inaccurate result or would produce errors in the methodology when they were to be processed. The filtered images that have been processed will be compared next to the the good quality images based on the tree surface area, enclosed surface area and horizontal center of area. The time-series for the 10-day analysis is shown in Figure 3.25.

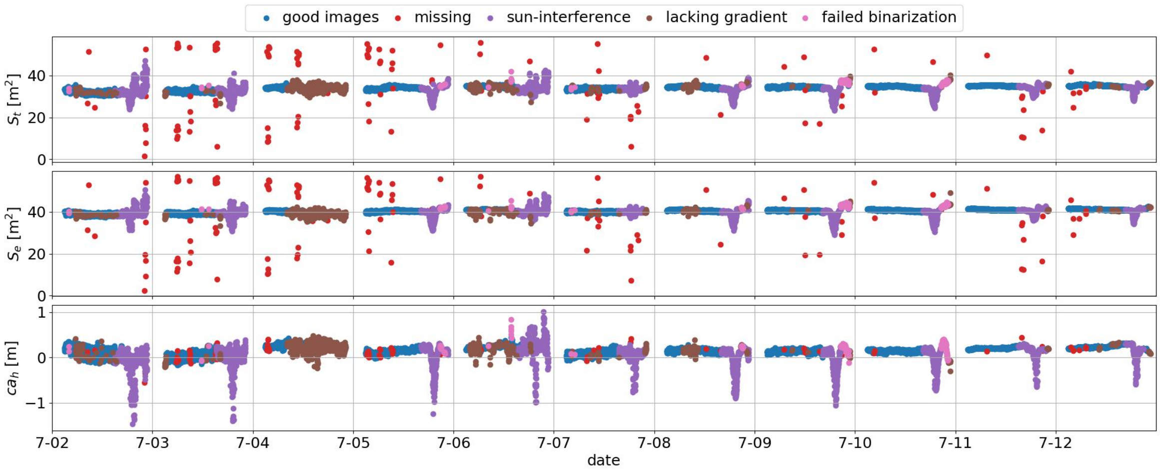

Figure 3.25: Time-series in 2019 for the 10-day analysis based on tree surface area, enclosed surface area and horizontal center of area where images that have been flagged as requiring to be filtered are indicated based on what they should have been filtered for.

This figure shows a much higher spread for the data, and many outliers can be identified. It is seen that missing data highly influences the area of the tree, both increasing and decreasing it. This is as the missing data is interpreted as being gray of colour. Based on whether the threshold will fall below or above this gray value, will decide whether the tree will highly increase or decrease in size. Furthermore, it is seen that images with missing data do not highly influence the horizontal center of area. This is expected as the missing data is not horizontally biased to a specific side to the image and simply cuts off the image at a certain height.

The sun interference is seen to decrease the surface area for later days in the time-series, but increase it for earlier days. A decrease will happen when the sun is at the edge of the tree and will create a lens flare which will cover the tree. This streak of brightness will be interpreted by the thresholding method as being part of the sky, decreasing the area of the tree. In the case of an increase in area, the brightness of the image might allow the automatic thresholding to find an optimal split above the value of clouds, adding parts of the clouds on the left side of the image to the area. Both these phenomena explain why almost always, the horizontal surface area moves to the left due to the sun. Either area is removed from the right side, or added to the left.

The lacking gradient is seen to not cause mayor changes in the surface area, but rather increase variance. This is due to the drops on the lens, occasionally counting as increasing the tree, and occasionally decreasing the tree. This extends itself to the horizontal center of area.

The failed binarization filters out over-estimations of the area. This is as the sky saturation is unequal, leading to parts of the sky to be included in the tree.

A closer look is taken at the images taken across two single days. The first of which is the third of July. This time-series is highlighted in Figure 3.26.

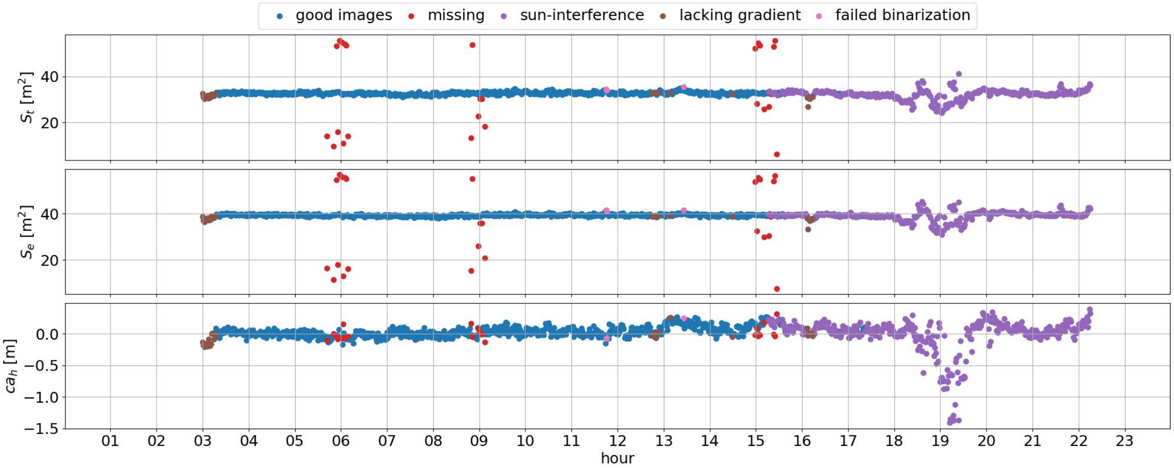

Figure 3.26: Time-series of the third of July, 2019, for the 10-day analysis based on tree surface area (top), enclosed surface area (middle) and horizontal center of area (bottom).

This day seems to have little interference except for the sun. It is seen that a large portion that has been classified as having sun interference might have been good quality images. This is an artifact of the way that the sun-interference separation has been trained. A high recall was sought after, which leads to a high sensitivity of classifying sun-interference. The images in the time-period have been investigated, and the reason that these are classified as having sun-interference is an exceptionally high case of fully saturated rolling clouds. From roughly 17:00 onward however, a lens-flare is present in the tree, creating a slowly increasing negative bias of roughly 50x50 pixels, which roughly translates to 0.25 m^2. This is almost a percent of the tree area. We see changes relative to the wind of roughly 1 percent per m/s, meaning this sensitivity is justified.

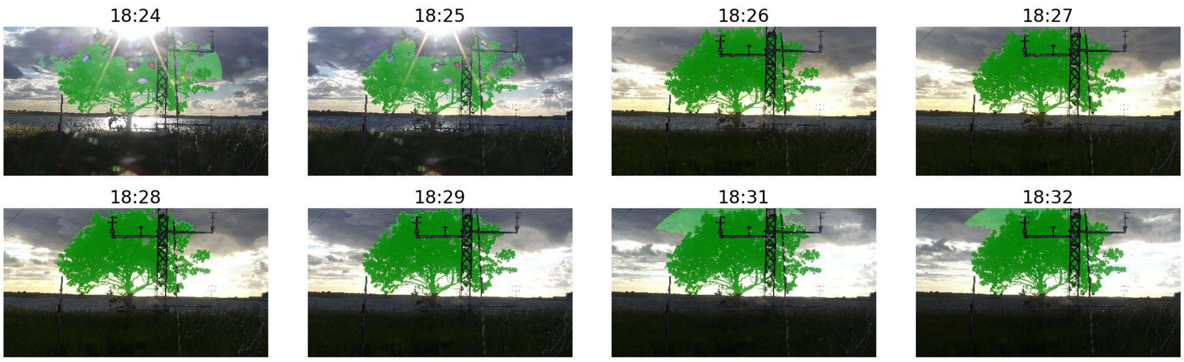

Figure 3.27: Visualisation of the decrease in surface area, which transfers to an increase due to the sun being blocked out by clouds.

Furthermore it is seen that roughly after 18:00, sun-interference is present in the forms of a strong area and horizontal center of area influence. First the area decreases and later, it increases. This increase is due to the sun being blocked out by a cloud, creating a bright light which makes part of the sky be included in the tree. This is seen in Figure 3.27.

A further highlight is made towards the fourth of July, where much rain is seen. This time series is seen in Figure 3.28.

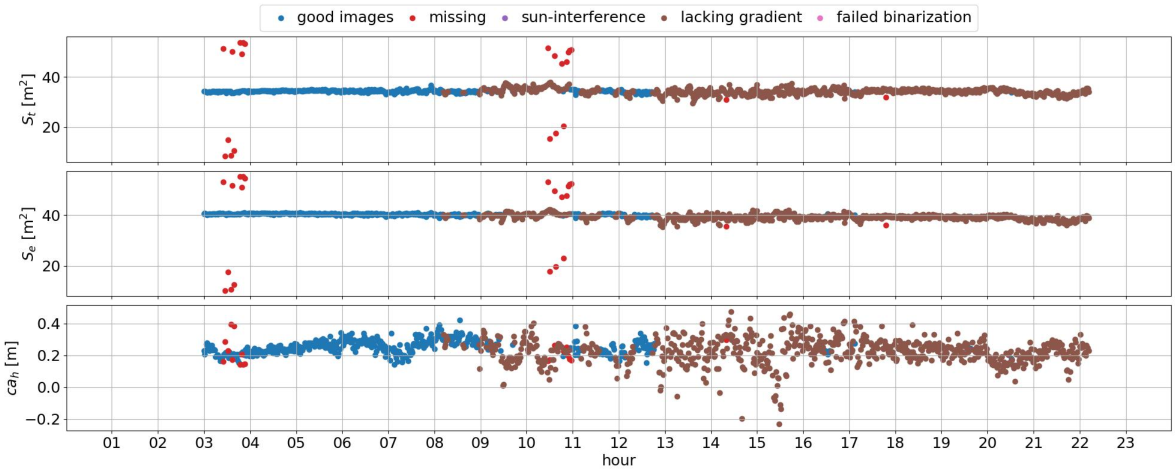

Figure 3.28: Time-series of the fourth of July, 2019, for the 10-day analysis based on tree surface area (top), enclosed surface area (middle) and horizontal center of area (bottom).

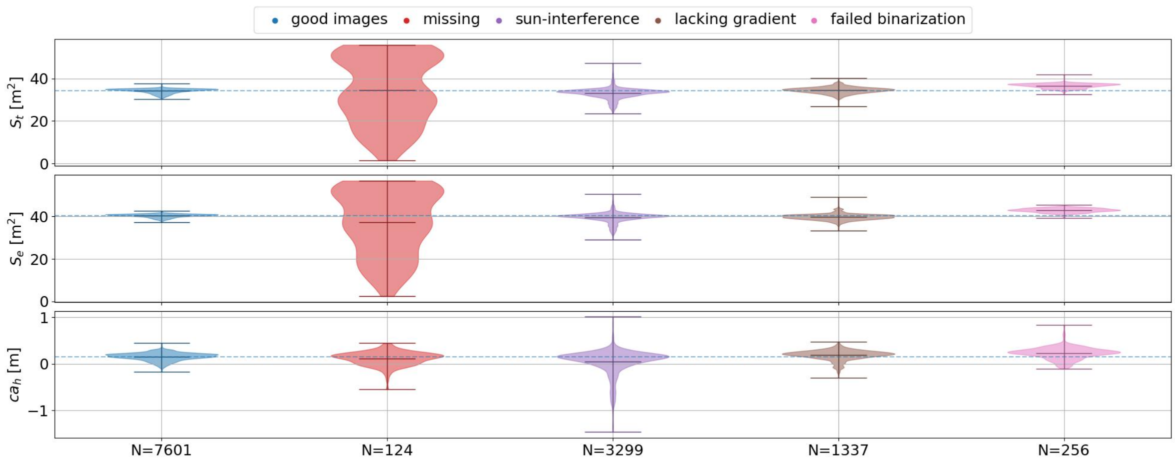

Figure 3.29: Violin plots for good quality and filtered images for the tree surface area (top) the enclosed surface area (middle) and the horizontal center of area (bottom).

From Figure 3.28, it is seen that the rain adds variance, as it sometimes adds and sometimes subtracts from the surface area. To get a better view of the distribution for each of the image aspects of each of the filtered out images, violin plots are shown in Figure 3.29

The missing images, are seen to have a large spread, and little is concluded from the shape of the distribution. This is partly due to the low amount of samples with missing data. A high spread in the surface areas is seen however. For the center of area, it is again seen that no noticeable bias or variation is induced for the horizontal surface area, as missing data does not discriminate horizontally.

The sun interference images does not show a clear bias in the surface areas, it does increase variance and includes some outliers. As the sun impacts in a somewhat similar way each day on the right side of the tree, it reduced area here and creates a negative bias for the horizontal center of area.

The lacking gradient images do not inhabit a bias, and are only seen to increase variance slightly, and include some lower outliers in surface area.

The failed binarization images are shown to not increase variance much for the surface area, but they do introduce a positive bias. This is because beyond just the tree, part of the sky is also classified as being part of the tree. This however, does not seem to have a directional bias, as there is no bias seen in the horizontal center of area. Here the variance is slightly higher, due to the area being pulled to either the left or the right.

To show the inhibiting factor of the bad quality images on investigating the parameters, they are plotted versus wind speed in Figure 3.30. It is indicated which data points would have been filtered out.

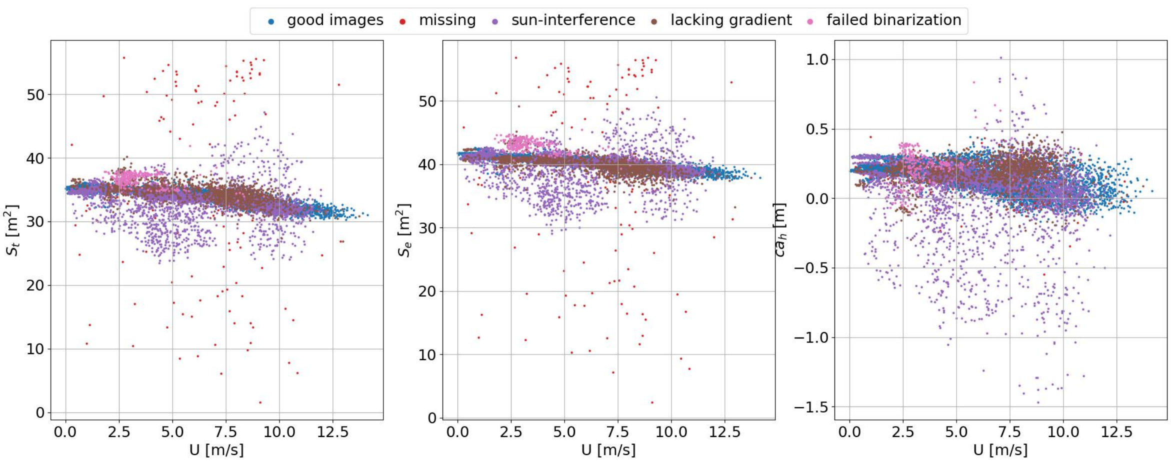

Figure 3.30: Correlation plot of tree surface area (left), enclosed surface area (middle) and horizontal center of area (right) for good quality images and images that are filtered out.

From this figure, it can be seen that an increase in spread and a bias is created by all of the faulty images, validating the importance of the filtering step in the analysis.

CHAPTER 4

Conclusion

Throughout this research, a detailed measurement period between the first of May 2019 until the 31^{st} of October 2019, on a single mature and relatively stiff tree has been analysed. This analysis has been done by relating wind statistics to image aspects retrieved from images made by a surveillance camera. These image aspects have been retrieved by means of a self-developed methodology which filters out unusable images and extracts relevant image aspects from left-over images.

The filtering process is accomplished by first automatically masking away pixels below the horizon, far away around the canopy and coinciding with an occluding meteorological mast. Secondly, 7 filters are applied. Empty images, night images, images with a camera misorientation, images with missing data within regions of interest, images with sun interference, images with rain interference and failed binarizations are all filtered out. During the process of filtering, 47% of images are excluded from analysis, noting that the sun interference caused by pointing the camera towards a region of the sky where the sun passes by, as one of the major contributions.

Each step in the filtering process has been optimized to reduce the unnecessary exclusion of images and increase the recall of eliminating images which would produce incorrect results. For filtering out images with sun interference, five methods have been compared with high ranges of parameters to find an optimal method of decreasing false positives while maintaining a recall higher than 98%. For the process of filtering out rain interference, a logistic model has been trained which selected hyper parameters based on two-layer cross validation, achieving an accuracy of roughly 94%.

As a method of validating the filtering process, images throughout a period of 10 days have been processed regardless of missing data, sun interference, lacking gradient or failed binarization. It was found that all filters had been justified as they introduced either an increase in variance in the data or introduced a bias.

Four image aspects have been extracted from each left-over image after the filtering process. These are the tree surface area, the enclosed surface area, the horizontal center of area and the vertical center of area. They are found by binarizing the tree and finding the circumference points. To find the location and of each pixel belonging to the tree, the position and orientation of the camera is computed. This allows for a physical representation of the surface area as well as the ability to locate the physical center of area relative to the base of the tree.

The binarization process uses thresholding of a single channel where the single channel has been optimized with three ground truth images to allow for the best separation between tree and non-tree pixels. The enclosed area is described by the area within the circumference and the tree area is described by the number of tree pixels found and adding an extrapolation of the porosity onto the occluding pixels. The surface area in pixels is converted into a physical area parameter by estimating the distance between the tree and the camera. It was noted that the definition of the enclosed area was dependent on the length scale used to find the circumference points, thus making it, and porosity calculated with it in general a less robust value.

The method has been validated by finding concurrence within expected behaviour of the tree and trends found in the data. These behaviours are the decrease in porosity and an increase in both tree surface area and enclosed surface area as the tree grows leaves. A further concurrence is a higher relative surface area reduction of the tree in a period with high amounts of leaves relative to periods with a low amount of leaves, as its general flexibility increases. Finally, the swaying direction of the horizontal center of area of the tree was found to be dependent on the wind direction and its magnitude to be dependent on the wind speed.

Unexpected and new findings from the analysis are an increase in surface area with wind speed for certain wind directions, indicating that the tree's ability to streamline in one direction comes at the cost of poor streamlining ability in another. This change is still evolutionarily beneficial as most trees will see high winds speeds mostly from one wind direction. Furthermore, as the tree is non-homogeneous, the sway of the tree is not necessarily in line with the wind direction, as an offset of 10 degrees with the line of sight of the camera has been found. Finally, the porosity was seen to initially decrease before increasing. This is as before the leaves start to streamline, increasing the porosity, the outline of the tree starts to streamline, decreasing the porosity.

Besides this 6-month analysis, an analysis of a higher temporal resolution on the period from the second of July 2019 until the 12^{th} of July 2019 is performed. The central focus of this study was on the effect of the surface area on the drag over wind speed relationship and the link between porosity and wind deficit.

The drag force is seen to not be proportional to the velocity squared as expected for a higher wind speed regime. The reasons for this have been identified as a change in drag coefficient and a change in area. The offset from 2 in exponent (Vogel exponent) between the drag and velocity was found to be -0.256 when assuming a stationary tree surface area. Considering a dynamic tree surface area, a value of -0.118 was found, concluding that slightly more than half of the offset in this exponent is due to the decrease in tree surface area, while other parts are to be explained by other physical phenomena. Finally, the wind deficit was found to not be fully linearly related to the porosity. This alludes to the possibility of skin friction drag playing a significant role in the wind deficit of a porous tree.

CHAPTER 5

Future Research

Some future research recommendations are presented.

Set-up

If a future experimental set-ups is created for a similar purpose four recommendations exist. In filtering, much data was filtered out to ensure no rain and sun interference. To ensure high availability of data, a rain guard can be attached to out-door camera's. A consideration towards the solar path across the sky can be made to not include it in the viewing direction of the camera. Thirdly, it is advised to avoid occlusion of the object of interest, as much uncertainty was added to the surface area by not knowing whether pixels behind the eastern meteorological mast belonged to the tree. Finally, some uncertainty existed in this research about the exact time at which the image was taken. If this is noted down precisely, a wind statistics of a smaller time interval centered around the image can be used.

Higher temporal resolution

A higher temporal resolution (around one second, or even higher) than 1-minute for the camera could be looked into. In this work, each image taken represents a snapshot of a large time-series of varying wind speeds causing a high addition of spread in the data. Only due to the large data size could conclusions be made in this case. Many influencing factors have been tried to be identified, but some are sure to have not been considered. By reducing the total period of investigation, these factors can be held constant with a higher confidence.

Remaining contribution of Vogel exponent

Roughly half of the Vogel exponent has been attributed to the shrinking frontal surface area of the tree. Hypotheses can be created towards the cause for the other contribution to the Vogel exponent. The inclusion of skin-friction drag for a high-speed regime is hinted at by a non-linear relationship between the wind deficit and the porosity. A follow-up study can be performed on the study by Vogel, 1989 which investigated the drag response of a hanging leaf with varying wind speeds in a wind tunnel. For this experiment Vogel had used a stationary surface area. A dynamic surface area can be used to record the transition from pressure drag to friction drag.

Top view camera

Non-homogeneous swaying was perceived in the tree. Much insight can be gained on the behaviour of the swaying of a single tree by mounting a camera above the tree and getting a bird's eye view. The swaying of the tree could be closely investigated. The tree could easily be identified by creating contrast on the ground near the tree.

APPENDIX A
Circumference point algorithm

Algorithm 1 Identify circumference points

```
Given: binary image I of shape (X, Y).
T ← [].
for x = 1 to X do
    for y = 1 to Y do
        if I(x, y) = 1 then
            T.Add((x, y))                          ▷ Add coordinates if pixel is active
        end if
    end for
end for
i_s ← 0                      ▷ Initially choose top-most left-most point as starting point
while True do
    startPoint ← T[i_s]
    nowPoint ← startPoint
    lookingDir ← 90                                ▷ Start looking north
    p ← []                                         ▷ Allocate p to store circumference points
    while True do
        p.Add(nowPoint)
        for i = 1 to 8 do
            step ← (Sign(cos(lookingDir)),Sign(sin(lookingDir)))
            checkPoint ← nowPoint + step
            if I(checkPoint) in T then
                Break
            else if i = 8 then
                Break out of while loop            ▷ A singular point, and nothing was found
            else
                lookingDir ← lookingDir +45        ▷ Look 45 degrees further
            end if
        end for
        nowpoint ← checkPoint                      ▷ Take a step to the found point
        lookingDir ← lookingDir −90                ▷ Look back to not skip right turns
        if nowPoint=startPoint then
            Break                                  ▷ Completed the circumference
        end if
    end while
    if Length(p) >= max(X, Y)/2 then
        Break                                      ▷ Big enough circumference found
    else
        i_s ← i_s + 1                              ▷ Retry with new starting point
    end if
end while
return p
```

APPENDIX B
High resolution filtering indication

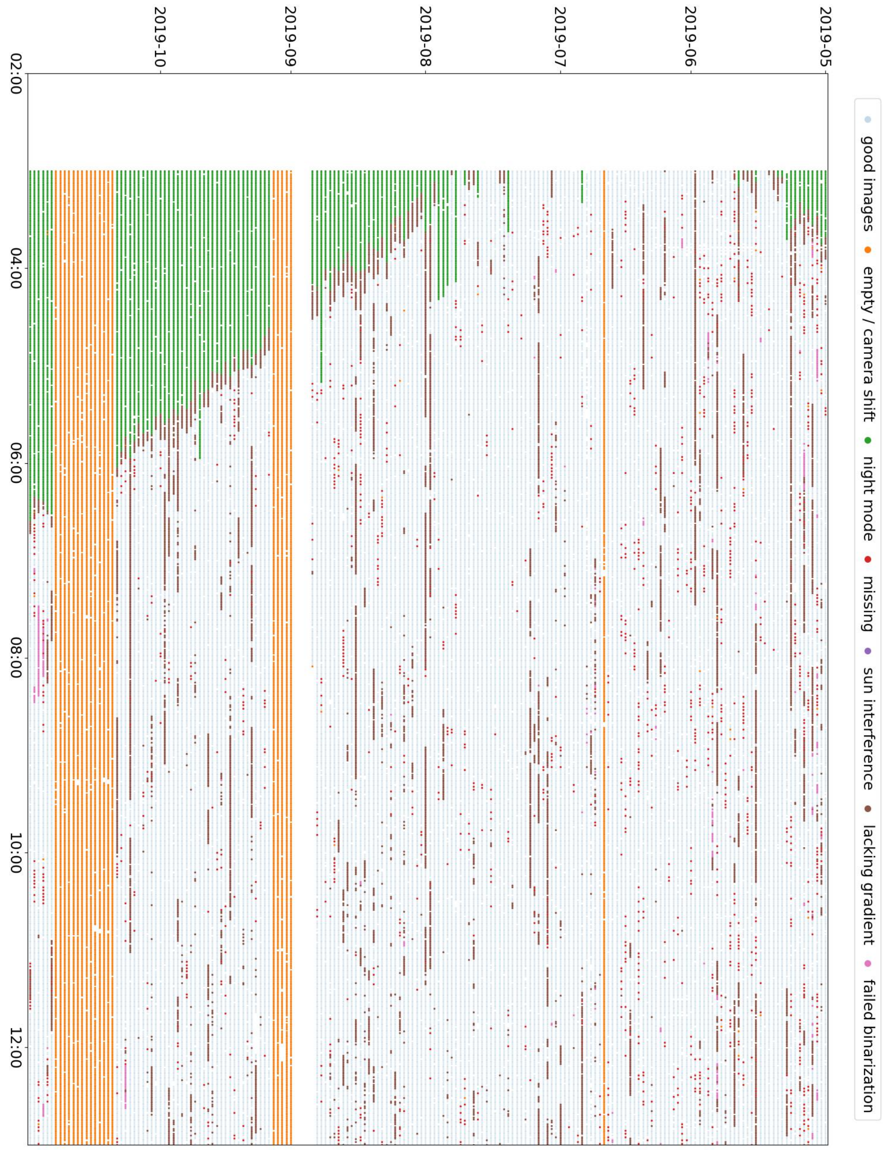

Figure B.1: Images filtered in time where the x-axis shows the day and the y-axis shows the time in the day. Hours from 2:00 until 13:00 are shown.

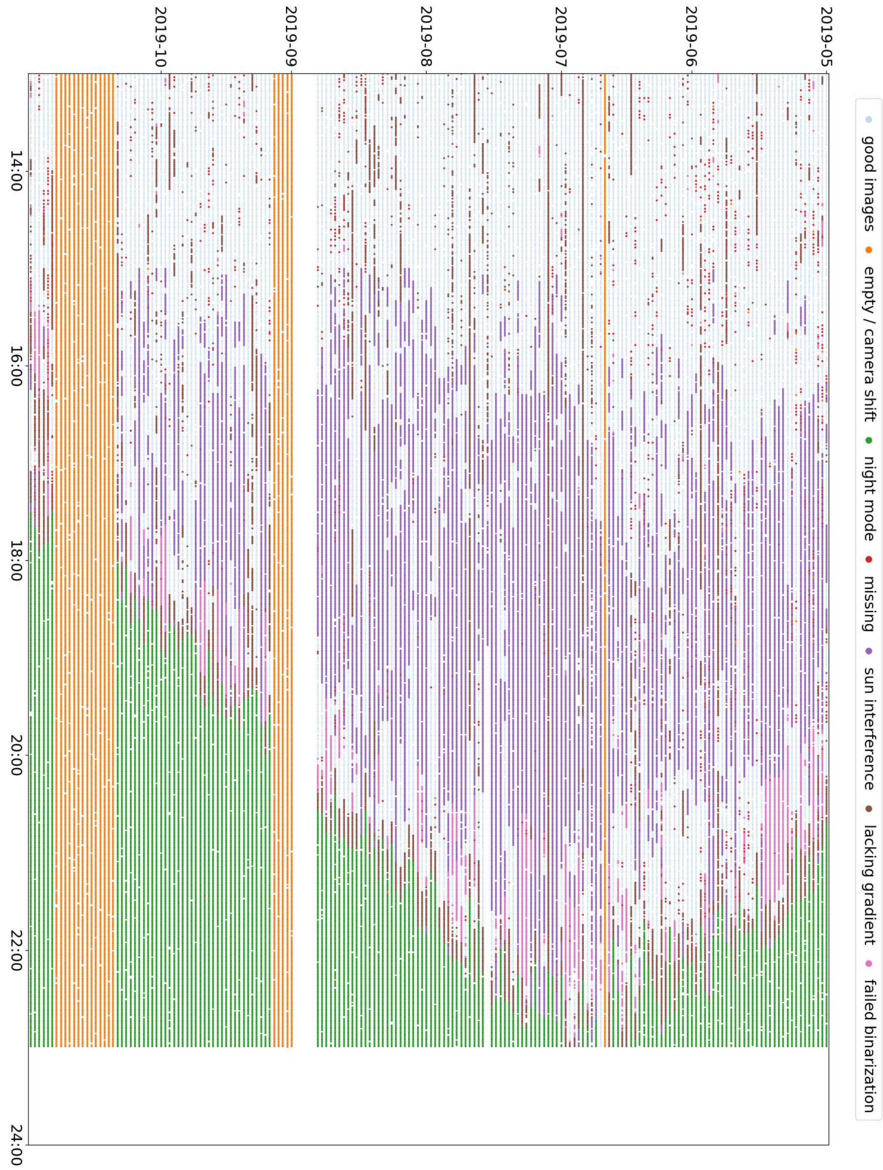

Figure B.2: Images filtered in time where the x-axis shows the day and the y-axis shows the time in the day. Hours from 13:00 until 24:00 are shown.

Bibliography

Agrawal, Rakesh and Ramakrishnan Srikant (2013). “Fast Algorithms For Mining Association Rules In Datamining”. In: *International Journal of Scientific & Technology Research* 2.12, pages 13–24.

Angelou, Nikolas, Ebba Dellwik, and Jakob Mann (2019). “Wind load estimation on an open-grown European oak tree”. In: *Forestry: An International Journal of Forest Research* 92.4, pages 381–392. ISSN: 0015-752X. DOI: 10.1093/forestry/cpz026.

Bache, D.H. (January 1986). “Momentum transfer to plant canopies: Influence of structure and variable drag”. In: *Atmospheric Environment (1967)* 20.7, pages 1369–1378. ISSN: 00046981. DOI: 10.1016/0004-6981(86)90007-7.

Boopathi, K et al. (2014). “Validation of accuracy of WAsP in vertical wind profile prediction”. In: *Journal of Renewable and Sustainable Energy*. Volume 6. 1, page 13143. DOI: 10.1063/1.4850255.

Cao, Jinxin, Yukio Tamura, and Akihito Yoshida (2012). “Wind tunnel study on aerodynamic characteristics of shrubby specimens of three tree species”. In: *Urban Forestry & Urban Greening* 11.4, pages 465–476. ISSN: 16188667. DOI: 10.1016/j.ufug.2012.05.003.

Cassiani, Massimo, G. G. Katul, and J. D. Albertson (March 2008). “The effects of canopy leaf area index on airflow across forest edges: Large-eddy simulation and analytical results”. In: *Boundary-Layer Meteorology* 126.3, pages 433–460. ISSN: 00068314. DOI: 10.1007/s10546-007-9242-1.

Dellwik, E. et al. (February 2019). “Observed and modeled near-wake flow behind a solitary tree”. In: *Agricultural and Forest Meteorology* 265, pages 78–87. ISSN: 01681923. DOI: 10.1016/j.agrformet.2018.10.015.

Dellwik, Ebba, Lars Landberg, and Niels Otto Jensen (2006). “WAsP in the forest”. In: *Wind Energy* 9.3, pages 211–218. ISSN: 10954244. DOI: 10.1002/we.155.

Denny, Mark and Brian Gaylord (2002). *The mechanics of wave-swept algae*.

Diener, Julien, Lionel Reveret, and Eugene Fiume (October 2006). “Hierarchical retargetting of 2D motion fields to the animation of 3D plant models”. In: *Computer Animation, Conference Proceedings*. Volume 02-04-Sept. Institute of Electrical and Electronics Engineers Inc., pages 187–195. ISBN: 3905673347.

Doaré, O, B. Moulia, and E De Langre (2004). “Effect of Plant Interaction on Wind-Induced Crop Motion”. In: *Journal of Biomechanical Engineering* 126.2, pages 146–151. ISSN: 01480731. DOI: 10.1115/1.1688773.

Driver, HE and AL Kroeber (1932). “Quantitative expression of cultural relationships”. In: *University of California Publications in American Archaeology and Ethnology* 81.4, pages 211–256.

Farquhar, T, J Z Wood, and J Van Beem (2000). "The kinematics of wheat struck by a wind gust". In: *Journal of Applied Mechanics, Transactions ASME* 67.3, pages 496–502. ISSN: 15289036. DOI: 10.1115/1.1311960.

Fix, Evelyn and J.L. jr. Hodges (1951). "Nonparametric Discrimination. 2. Small-sample performance in normal populations." In: *The Annals of Mathematical Statistics* 22.3.

Garnier, Jean Guillaume and Adolphe Quételet (1838). *Correspondance mathématique et physique, Volume 10*. Impr. d'H. Vandekerckhove.

Gosselin, Frédérick P. (2019). *Mechanics of a plant in fluid flow*. DOI: 10.1093/jxb/erz288. arXiv: 1905.08055.

GWEC (2019). *Global Wind Report 2018*. Technical report. Global Wind Energy Counci, page 61.

Hoerner, S.F. (1965). "Fluid-dynamic drag: practical information on aerodynamic drag and hydrodynamic resistance". In: *Journal of Physics: Conference Series*. Volume 75. By the Author, page 450. ISBN: 9991194444. DOI: 10.1088/1742-6596/75/1/012007.

Högström, Ulf and Ann Sofi Smedman (2004). "Accuracy of sonic anemometers: Laminar wind-tunnel calibrations compared to atmospheric in situ calibrations against a reference instrument". In: *Boundary-Layer Meteorology* 111.1, pages 33–54. ISSN: 00068314. DOI: 10.1023/B:BOUN.0000011000.05248.47.

Kane, Brian and E. Thomas Smiley (August 2006). "Drag coefficients and crown area estimation of red maple". In: *Canadian Journal of Forest Research* 36.8, pages 1951–1958. ISSN: 00455067. DOI: 10.1139/X06-086.

Kapur, J. N., P. K. Sahoo, and A. K.C. Wong (1985). "A new method for gray-level picture thresholding using the entropy of the histogram." In: *Computer Vision, Graphics, & Image Processing* 29.3, pages 273–285. ISSN: 0734189X. DOI: 10.1016/0734-189X(85)90125-2.

Kittler, J. and J. Illingworth (1986). "Minimum error thresholding". In: *Pattern Recognition* 19.1, pages 41–47. ISSN: 00313203. DOI: 10.1016/0031-3203(86)90030-0.

Koizumi, Akio et al. (June 2010). "Evaluation of drag coefficients of poplar-tree crowns by a field test method". In: *Journal of Wood Science* 56.3, pages 189–193. ISSN: 14350211. DOI: 10.1007/s10086-009-1091-8.

Kolmogorov, A N (1941). "The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds' Numbers". In: *Doklady Akademiia Nauk SSSR* 30, pages 301–304. ISSN: 1364-5021. DOI: 10.1098/rspa.1991.0075.

Lee, Y. J., R. I. Alfaro, and G. A. Van Sickle (1983). "Tree-crown defoliation measurement from digitized photographs." In: *Canadian Journal of Forest Research* 13.5, pages 956–961. ISSN: 00455067. DOI: 10.1139/x83-127.

Levenberg, Kenneth (1944). "A method for the solution of certain non-linear problems in least squares". In: *Quarterly of Applied Mathematics* 2.2, pages 164–168. ISSN: 0033-569X. DOI: 10.1090/qam/10666.

Li, Zhenjia, J. D. Lin, and D. R. Miller (April 1990). "Air flow over and through a forest edge: A steady-state numerical simulation". In: *Boundary-Layer Meteorology* 51.1-2, pages 179–197. ISSN: 00068314. DOI: 10.1007/BF00120467.

Liu, Dong C. and Jorge Nocedal (October 1989). "On the limited memory BFGS method for large scale optimization". In: *Mathematical Programming* 45.1-3, pages 503–528. ISSN: 00255610. DOI: 10.1007/BF01589116.

Mandelbrot, Benoit (1967). "How long is the coast of Britain? Statistical self-similarity and fractional dimension". In: *Science* 156.3775, pages 636–638. ISSN: 00368075. DOI: 10.1126/science.156.3775.636.

Mitchell, Tom M. (Tom Michael) (1997). *Machine Learning*. Illustrate. McGraw-Hill Education, page 414. ISBN: 978-0-07-042807-2.

Mizoue, Nobuya and Akio Inoue (2001). "Automatic Thresholding of Tree Crown Images". In: *Journal of Forest Planning* 6.2, pages 75–80. ISSN: 1341-562X. DOI: 10.20659/jfp.6.2_75.

N., Otsu; (1979). "A threshold selection method from gray-level histograms." In: *IEEE Trans. Syst. Man. Cybern.* 9, pages 62–66.

Nelder, J. A. and R. Mead (1965). "A Simplex Method for Function Minimization". In: *The Computer Journal* 7.4, pages 308–313. ISSN: 0010-4620. DOI: 10.1093/comjnl/7.4.308.

Peña, Alfredo, Ebba Dellwik, and Jakob Mann (2019). "A method to assess the accuracy of sonic anemometer measurements". In: *Atmospheric Measurement Techniques* 12.1, pages 237–252. ISSN: 18678548. DOI: 10.5194/amt-12-237-2019.

Py, Charlotte, Emmanuel De Langre E., and Bruno Moulia (2006). "A frequency lock-in mechanism in the interaction between wind and crop canopies". In: *Journal of Fluid Mechanics* 568, pages 425–449. ISSN: 00221120. DOI: 10.1017/S0022112006002667.

Py, Charlotte, Emmanuel De Langre, et al. (2005). "Measurement of wind-induced motion of crop canopies from digital video images". In: *Agricultural and Forest Meteorology* 130.3-4, pages 223–236. ISSN: 01681923. DOI: 10.1016/j.agrformet.2005.03.008.

Rathmann, Ole et al. (1996). "Assessing the accuracy of WAsP in non-simple terrain". In: *BWEA18 conference, 24-27 September 1996, University of Exeter, Exeter, United Kingdom*, pages 6–11.

Rosenblatt, Murray (1956). "Remarks on Some Nonparametric Estimates of a Density Function". In: *The Annals of Mathematical Statistics* 27.3, pages 832–837. ISSN: 0003-4851. DOI: 10.1214/aoms/1177728190.

Rudnicki, Mark, Stephen J Mitchell, and Michael D Novak (2004). "Wind tunnel measurements of crown streamlining for drag relationships for three conifer species". In: *Canadian Journal of Forest Research* 34.3, pages 666–676. ISSN: 00455067. DOI: 10.1139/x03-233.

Samuel, A. L. (1959). "Some studies in machine learning using the game of checkers". In: *IBM Journal of Research and Development* 3.3, pages 210–229. ISSN: 00188646. DOI: 10.1147/rd.441.0206.

Smith, Steven W. (1997). "Moving Average Filters". In: *THE SCIENTIST & ENGINEER'S GUIDE TO DIGITAL SIGNAL PROCESSING*. 1st. California Technical Publishing. Chapter 15, pages 277–284. ISBN: 0966017633.

Tan, Pang-Ning et al. (2019). *Introduction to Data Mining*. 2nd. Pearson. ISBN: 0133128903.

Troen, Ib and Erik Lundtang Petersen (1989). "European wind atlas - Errata". In: page 656.

Turing, A.M. (1950). “Computing Machinery and Intelligence”. In: *Mind a quarterly review of psychology and philosophy* 59.236, pages 433–460.

Vogel, Steven (February 1984). “Drag and Flexibility in Sessile Organisms”. In: *American Zoologist* 24.1, pages 37–44. ISSN: 0003-1569. DOI: 10.1093/icb/24.1.37.

— (1989). “Drag and reconfiguration of broad leaves in high winds”. In: *Journal of Experimental Botany* 40.8, pages 941–948. ISSN: 00220957. DOI: 10.1093/jxb/40.8.941.

Vollsinger, Stephan et al. (2005). “Wind tunnel measurements of crown streamlining and drag relationships for several hardwood species”. In: *Canadian Journal of Forest Research* 35.5, pages 1238–1249. ISSN: 00455067. DOI: 10.1139/x05-051.

Weibull, W. (1951). “A Statistical Distribution Function of Wide Applicability”. In: *Journal of Applied Mechanics* 13.3, pages 293–297.