University of Dublin

TRINITY COLLEGE

Weather Analysis using Computer Vision

David Young
B.A.(Mod.) Business and Computing
Final Year Project April 2014
Supervisor: Dr Kenneth Dawson-Howe

School of Computer Science and Statistics

O'Reilly Institute, Trinity College, Dublin 2, Ireland

DECLARATION

I hereby declare that this project is entirely my own work and that it has not been submitted as an exercise for a degree at this or any other university

_____ _____

Name Date

Acknowledgements

I wish to thank Kenneth Dawson-Howe for his valuable insight and dialogue throughout the project work.

I also wish to thank my family and friends for their constant support during the year.

Abstract

Analysis of weather is an important emerging topic in Computer Vision. This work presents usage of Computer Vision techniques for analysis of rain, light levels and wind.

Analysis of rain is based on direct observation of moving raindrops allowing estimation of rain intensity. Light levels are inferred by configuring the camera to act as a light meter. Wind is studied based on the movement of tree features with estimation of windspeed from such movement.

Experimental observations are correlated with theoretical analysis or independent measurements.

Contents

Acknowledgements i

Abstract ii

Chapter 1 – Introduction 1

1.1. Aim 1

1.2. Motivation – Deployment Opportunities 1

1.3. Computer Vision for Analysis of Weather 1

1.4. Report Structure 2

Chapter 2 – Literature Review 3

2.1. Scope, Objective 3

2.2. Rain Analysis 3

2.2.1. Fundamental Properties of Rainfall and Effects relevant to CV Techniques 3

2.2.2. Local Climatic Information 5

2.3. Light Analysis 6

2.4. Wind Analysis 6

2.4.1. Inferring Wind Characteristics from Movement of Tree Features 7

2.4.2. Wind Indicators 8

2.4.3. Local Wind Characteristics 8

2.5. Summary 8

Chapter 3 – Experimental Setup 9

3.1. Introduction 9

3.2. Location Characteristics and Field of View 9

3.3. Camera 10

3.4. Calibration Equipment 11

3.4.1. Rain 11

3.4.2. Light 11

3.4.3. Wind....12
3.5. Summary....12
Chapter 4 – Rain....13
4.1. Objective....13
4.2. Effects of Rain on Images....13
4.3. Experimental Conditions and Theoretical Analysis....13
4.4. Experimental Setup....16
4.5. Process....16
4.6. Results from Analyses of Video Recordings....20
4.7. Analysis....21
4.7.1. Strengths....21
4.7.2. Limitations....22
4.7.3. Future Work....22
4.8. Conclusions....22
Chapter 5 – Light....23
5.1. Objective....23
5.2. Effect of Light on Images....23
5.3. Range of Light Intensity....23
5.4. Light Effects....25
5.5. Approach....25
5.6. Results....26
5.7. Analysis....28
5.7.1. Success....28
5.7.2. Limitations....28
5.7.3. Future Work....28
5.8. Conclusions....28

Chapter 6 – Wind 29
6.1. Objectives 29
6.2. Effect of Wind on Images 29
6.3. Wind Characteristics 30
6.3.1. Initial Analysis 30
6.3.2. Central Movement 30
6.3.3. Peripheral Movement 31
6.4. Approach 31
6.5. Results 36
6.6. Analysis 37
6.6.1. Successes 37
6.6.2. Limitations 38
6.6.3. Future Work 38
6.7. Conclusions 38
Chapter 7 – Conclusions 40
7.1. Overview 40
Appendix 1. Light-Meter Data 41
Appendix 2. Anemometer Data 41
Appendix 3. Detailed Analysis of Rainfall Characteristics 41
References 43
Attached DVD 45

Chapter 1 – Introduction

1.1. Aim

The aim of this work is to use Computer Vision techniques to characterise key aspects of weather, principally of rainfall and of wind, exploring opportunities and limitations.

1.2. Motivation – Deployment Opportunities

Using Computer Vision techniques for analysis of weather can offer valuable functionality. This is particularly the case if readily-available cameras can make use of such techniques.

Weather analysis and forecasting is currently undertaken without detailed study of particular areas that may have "microclimates". Aggregation of weather data from a network of cameras could aid understanding and short-term, precise prediction of weather in localised areas.

Weather analysis techniques are also important to vehicles and to traffic management. Automatic activation of traffic warning signs, including speed limit reduction, upon sensing of severe levels of rain and/or wind may for example be facilitated by usage of the techniques discussed here.

1.3. Computer Vision for Analysis of Weather

The principal weather characteristics that are likely to be of interest include:

- Precipitation – in most cases rain (in the relevant geographical area)
- Light level – as a proxy for whether the day is "gloomy" or bright.
- Wind
- Temperature

Of these, temperature is readily evaluated using electronic thermal sensors, which may be incorporated in any event within the camera system. This is not a topic for usage of "Computer Vision" techniques with standard cameras, and thus is not considered further.

Computer Vision techniques are by definition limited to what the camera can observe within its field of view, subject of course to any controlled pan/tilt/zoom operations that may be available.

In the case of rain, observable effects include:

- Raindrop traces visible on images
- Colour changes on some surfaces, such as concrete
- Splashing in “puddles”, typically formed after rain accumulation, or in ponds

Of these, the first approach is likely to be preferred. This has the advantage of giving real-time information whereas the other approaches give historical or “integrated” metrics of rain, as well as being quite imprecise and lacking robust performance in the context of variation in the other principal weather aspects such as wind and light characteristics.

In the case of light, an effective approach is to take the overall light level within the image. If this can be correlated to the overall “lux” figure, then this is likely to give the required indication.

In the case of wind, the observable effects are likely to be related either to “windsock-type” elements – rarely found in most scenes – or else by analysing the motion of vegetation, principally tree branches and leaves. The main challenge here is that response of trees to wind is little understood, and the movement may be associated more with wind fluctuations rather than correlated with average windspeed.

1.4. Report Structure

The work has been focused on analysing these phenomena. Chapter 2 provides a review of literature in this area, chapter 3 describes the experimental setup, in terms of the field of view and equipment deployed. Chapters 4, 5 and 6 focus in more detail on analysis respectively of rain, light and wind, with overall conclusions presented in chapter 7.

Chapter 2 – Literature Review

2.1. Scope, Objective

The principal focus of this work is on usage of Computer Vision ("CV") techniques to infer rainfall and wind characteristics from observing a scene. Capture of light conditions is also likely to be readily feasible by usage of camera properties. A review of literature is thus presented under these headings.

2.2. Rain Analysis

2.2.1. Fundamental Properties of Rainfall and Effects relevant to CV Techniques

In analysis of rainfall one reviews initially the literature on usage of relevant CV techniques for evaluation of rain. The physical properties of raindrops and rainfall data also need to be studied. This provides a theoretical framework within which experimental results can be analysed and which will allow interpretation of data as collected.

Studies of the effects of rain on images have been documented by Kshitiz Garg and Shree Nayar at Columbia University (Garg and Nayer (2007)). Their work has been further developed by Brewer and Liu (2008). The principal goal of these works was associated with removal of rain-induced "streaking" on video sequences as in broadcasting of outdoor events.

Raindrop properties need to be understood, principally the size distribution and the velocity. The raindrop distribution by size corresponding to given rainfall rates has been modelled by successive authors including Marshall and Palmer (1948), and an overview of distribution modelling has been included in the work by Uijlenhoet (2001). The Marshall and Palmer result is as shown in Fig. 2.1(a), where "D" is the drop diameter, seen to have a typical size between 1mm and 2mm for moderate rain levels, and N(D) corresponds to the drop distribution with size D.

A raindrop is essentially a small mass of water held together by surface-tension effects. This surface tension effect breaks down as the volume/(surface area) ratio increases, such that

raindrops above approximately 5mm in diameter become increasingly unstable and typically break into smaller drops.

Falling raindrops in air reach a terminal velocity given by the balance between gravitational forces and air resistance, with this terminal velocity given by

$$v = 200\sqrt{a}$$

with v the velocity and a the radius of the raindrop, in SI units (as for example cited in Garg and Nayer (2007)).

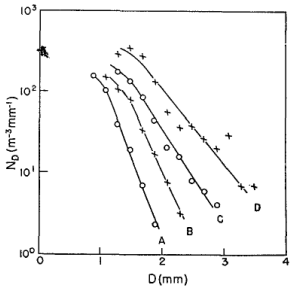

FIG. 1. Distribution of number versus diameter for raindrops recorded at Ottawa, summer 1946. Curve A is for rate of rainfall 1.0 mm hr^{-1}, curves B, C, D, for 2.8, 6.3, 23.0 mm hr^{-1}. $N_D\delta D$ is the number of drops per cubic meter, of diameter between D and $D + \delta D$ mm. Multiplication by 10^{-4} will convert N_D to the units of equation (2).

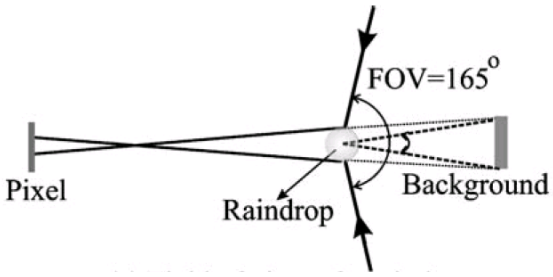

Figure 2.1(a) – left – showing the Marshall-Palmer Distribution estimates for raindrop size at differing rainfall rates, Figure 2.1(b) – right – showing the field of view due to the refraction effect within a raindrop

Raindrops are shown to be very effective at refracting incident light across a wide angle, including from the sky, (figure 2.1(b)), with minor contributions to radiance coming from reflection effects. This property makes raindrops appear materially brighter than the background image sections present in almost any field of view.

A camera typically has an exposure time of approximately 25ms-40ms. An object that is in motion during this exposure time will correspond to a "trail" being seen on the corresponding image, and given the enhanced brightness of raindrops they will appear as bright "streaks".

An image sequence containing falling raindrops may be processed to highlight these "streaks". A subtraction process is used, where the luminance component of the immediately preceding image is subtracted from the corresponding component in the "new" image. This results in the "streaks" becoming particularly apparent, as in figure

2.2(a). If detailed analysis of these streaks is needed, then a binary processed image may then be derived using a "threshold" approach.

A relevant concept introduced is that of the Rain Visible Region. This is determined by the ability of the camera to identify individual raindrops, with the visible region dependent on drop size and camera resolution. The plot of figure 2.2(b) illustrates this concept.

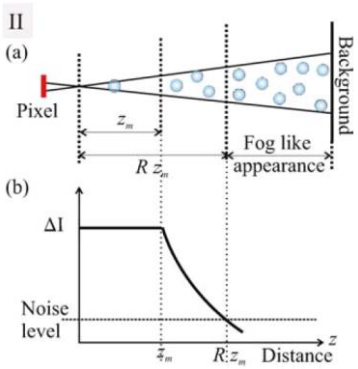

Figure 2.2(a) – left – showing streaks visible in a processed image. Figure 2.2(b) – right – showing the Rain Visible Region Rz_m – as defined by Garg and Nayer (2007).

The concept of "Rain Visible Region" is important in relating the number of streak events to rainfall levels, and needs to be complemented by information as to the depth of field of the camera used. This is central to provision of the theoretical background accompanying the experimental work presented in Chapter 4, and will be discussed further in that section.

Whilst the additional work in Garg and Nayer (2007) (and also in Brewer and Liu (2008)) is focused on "destreaking" for broadcast and motion picture applications, there is mention of the "Camera-Based Rain Gauge". This was able to make use of a camera with a suitable depth of field, and the aspect ratio (length/width of streak) was used to "qualify" streaks as being associated with raindrops.

2.2.2. Local Climatic Information

It is necessary to be cognisant of the range of weather conditions likely to be seen at the experimental location in Dublin, Ireland.

Climatic conditions for the Dublin area indicate an average rainfall of 750mm-1000mm annually, with average hourly rainfall levels typically being 1mm/hr to 2mm/hr (Met Eireann, 2014). Peak levels are often estimated as closer to 6mm/hr.

There is minimal Irish historical data available for instantaneous rainfall levels, largely due to the usage of integrating measurement approaches which capture rainfall amounts over periods of hours or days, although the move to radar data collection is permitting real-time information to be collated.

Raindrop size distribution data for Ireland – comparable to the Marshall-Palmer data – does not appear to be available. This distribution pattern cited has however been found relevant (e.g. in studies reported in Uijlenhoet (2001)) for most temperate regions and is thus used in background analyses.

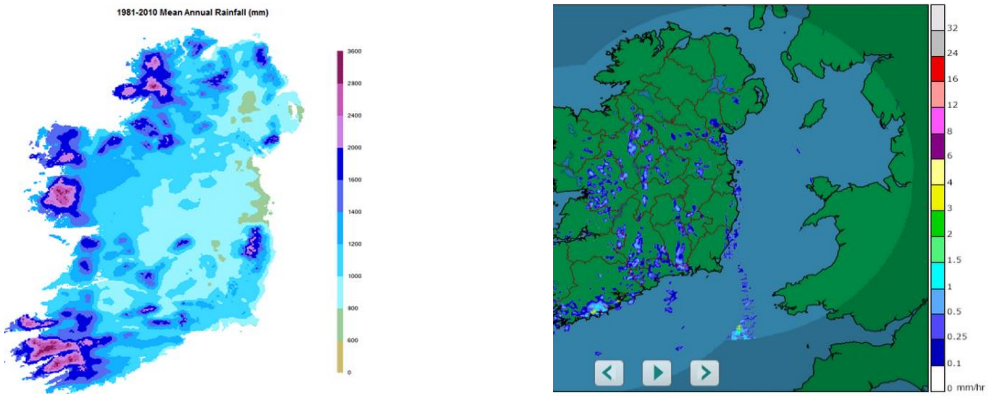

Figure 2.3(a) – left – showing annual rainfall characteristics, Figure 2.3(b) – right – showing rainfall radar information as presented in www.met.ie

2.3. Light Analysis

Whilst Computer Vision techniques have been previously used for sky analysis and cloud tracking, the selected location and the preference for a single camera with its limited internal Pan-Tilt-Zoom characteristics suggested that the main metric available in this context related to the overall light level.

The approach here essentially involves calibration of the camera for use as a photometer (light-meter). Such a calibration has been reported in several photographic industry publications, with an industrial deployment of this technique being cited in Cattoen et al (2005).

2.4. Wind Analysis

It is possible for Computer Vision techniques to be used to observe flag or windsock movement or to see movement of leaves and branches of trees.

If a flag or windsock were available within the field of view, then observation of its elevation and direction could allow some wind characteristics to be inferred. In the absence of such a feature, it is necessary to rely on interpreting tree movement to infer wind characteristics.

2.4.1. Inferring Wind Characteristics from Movement of Tree Features

One approach to this task has been presented in Lu and Ohya (2004). This uses a learning algorithm, where historical data as to movement of key features of a tree as a function of incident wind strength and direction is accumulated. Given that realistic learning could take some time and that either an aerial camera or a number of near-ground-based cameras would be required, this approach was considered impractical.

A more fundamental approach could use studies of "Tree Dynamics", and this is an area that has received increasing attention in recent years. Important contributions in this area have come from James et al (2006) and from authors at the 2009 Freiberg Conference (Mayer and Schindler (Eds) 2009). A primary motivation for such work has come from the fact that tree movement is clearly a commonplace event that is poorly understood, and study in this area is aligned with a general move to "precision agriculture". There is also an "application-oriented" motive in that understanding tree dynamics can allow more realistic evaluation of safety issues associated with trees in storms.

Modelling of trees has proceeded from the very simplistic (figure 2.4(a)) through to more complex variants as shown in figures 2.4(b) and 2.4(c).

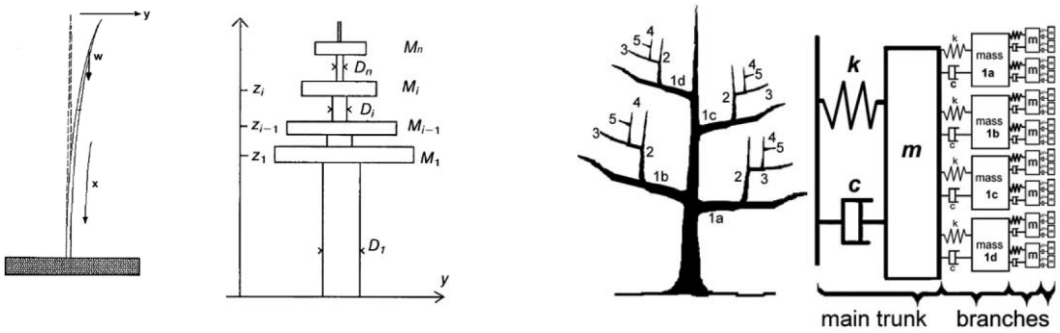

Figure 2.4(a) – left, Figure 2.4(b) – centre – and Figure 2.4(c) – right – showing increasing complexity in modelling of trees – (James et al, 2006)

A "complex" tree is therefore likely to be one with several "levels" of branching, each of which will have individual mechanical mass-spring-damper characteristics.

2.4.2. Wind Indicators

The consequence of this analysis is that movement at the periphery of a tree will be very complex, being quasi-random in nature and likely to be dependent on time derivatives of wind rather than solely on windspeed. A further factor is that trees are usually planted in groups or "stands", rather than as isolated trees.

In terms of getting an indicator of windspeed, with naturally a sensitivity to wind direction, it is possible to look at the movement of some feature closer to the centre of the tree, preferably an isolated tree. Here the dynamic effects are simpler and there is a realistic expectation – to be verified in chapter 6 – that the tilt of a tree may be correlated with windspeed. The challenge of course is that movement of large trees in wind will be minimal, and as a result it is necessary to select a small tree whose tilt will correspond to a meaningful displacement in terms of pixels within the field of view.

2.4.3. Local Wind Characteristics

Wind characteristics for Ireland have been documented by Met Eireann. The mean average windspeed for the Dublin area is close to the 6m/s contour, corresponding to approximately 20km/hr. Wind characteristics will be highly dependent on local topographical features, and this is discussed in subsequent treatment.

2.5. Summary

This review of prior work indicates that there is some relevant material in the area of deployment of Computer Vision techniques for rain analysis which can inform the approach to be taken. Allowing usage of cameras with less precise depth-of-field resolution is important, as is ensuring robustness in windy conditions, and these can constitute topics for chapter 4.

In the case of light measurement, being able to use a camera as a light-meter is a technique that has been used and whose validity will be evaluated in chapter 5.

Studies of wind effects relevant to deployment of CV techniques have largely been irrelevant to the conditions obtaining in this work, with further opportunity for original work in this context, discussed in chapter 6.

Chapter 3 – Experimental Setup

3.1. Introduction

This chapter outlines characteristics of the location chosen for usage of the Computer Vision techniques, and describes equipment used for calibration.

3.2. Location Characteristics and Field of View

The area chosen for all experimental work was in a suburban garden, approximately 8km south of the Trinity College Dublin main campus, as shown in figure 3.1(a) and 3.1(b).

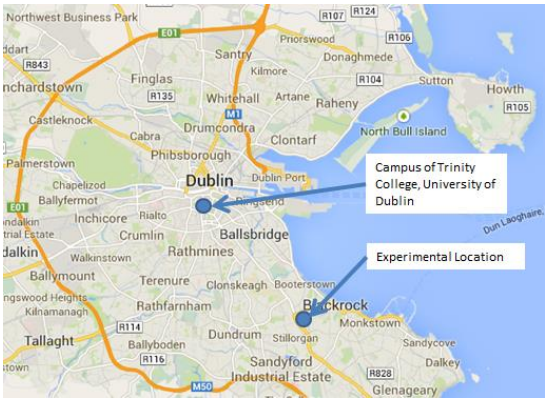

Figure 3.1(a) – left – showing the location relative to Dublin city and figure 3.1(b)– right – in additional detail, at 53°17’46”N and 6°12’8”W.

The garden is relatively sheltered by trees or buildings on three sides. This is illustrated in Figure 3.1(b) above.

Two Field of View (“FOV”) settings were used. For analysis of rain and wind, a small FOV was used, corresponding to a horizontal camera angle of 11.5 degrees. For analysis of light, a much larger FOV was used, with a horizontal camera angle of 50 degrees. The two FOVs are illustrated in Figure 3.2.

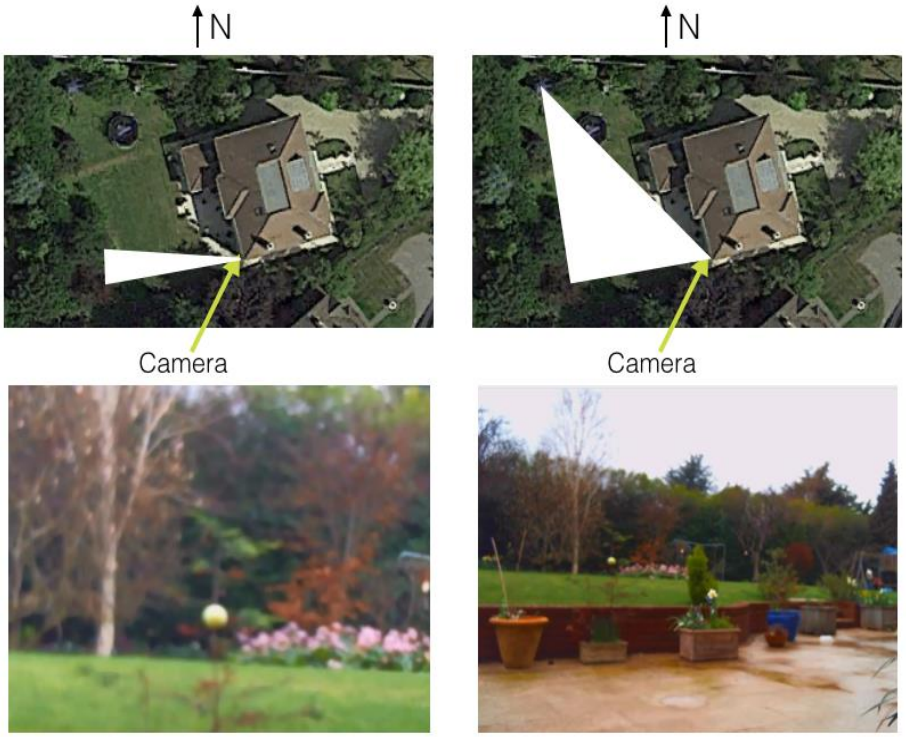

Figure 3.2(a) – Top Left – Showing the field of view (approximately 11.5°) used in wind and rain analyses, figure 3.2(b) – Top Right – showing the wider field of view (approximately 50°) as used in light analysis, with figure 3.2(c) and 3.2(d) showing the respective field-of-view images corresponding to the upper figures.

3.3. Camera

The camera used for all of the observational work was the Logitech C615 HD Webcam (Logitech, 2014). The C615 is programmable, with an associated API allowing control of key parameters such as Pan/Tilt/Zoom, Focus and Gain Control (QuickCam Team, 2007). This allowed the camera settings to be adjusted under software control to optimise parameters for the phase of the project being undertaken.

This functionality was enhanced with additional code as needed for work in this project. The code allowed optimal setup of camera parameters on initialisation and as required during the video sequences.

The camera was located on a windowsill approximately 0.7m from the ground. It was housed in a plastic container to protect it from rain. The container had a cutout for the image window and the camera was positioned centrally. The container was also provided

with holes to allow drainage of water from any incidental rain. The camera and box were secured in position, to ensure consistent results, using sealing compound.

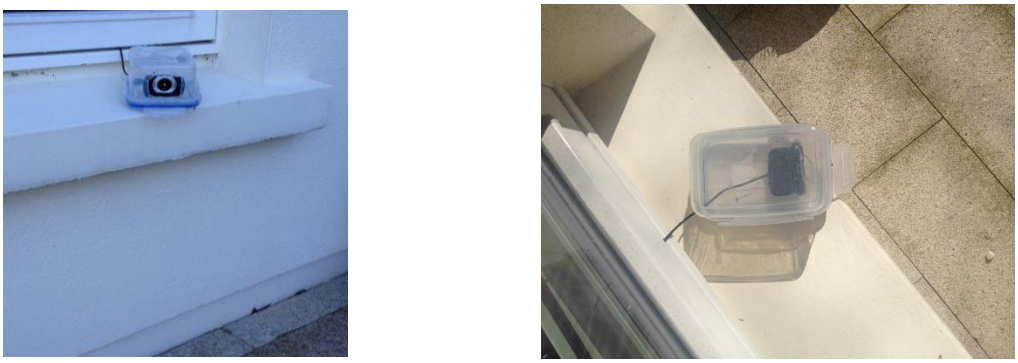

Figure 3.3 Camera positioning

3.4. Calibration Equipment

Additional equipment was used to calibrate the measurement process for relevant parameters which are highlighted below.

3.4.1. Rain

Calibration of rainfall could conceivably involve usage of a commercial rain-gauge, and indeed one was available as part of the system including the anemometer. However, this was judged to be impractical in that the rain gauge accumulates rain over a considerable period of time, and does not offer the capability to measure rain over the short intervals needed for calibration of a real-time measurement system as present here. The theoretical analysis as undertaken is thus important in confirming consistency of readings.

3.4.2. Light

A commercial light-meter, normally intended for photographic applications, was used for calibration of the camera for measurement of light conditions, as will be treated in chapter 5. The light-meter in question was a Maplin N76CC, with specifications available in the appendix.

Figure 3.3(a) – left – showing the light-meter, and Figure 3.3(b) – right – showing the anemometer as used.

3.4.3. Wind

The wind measurements could be independently verified using a commercial anemometer, as shown in figure 3.3(b). This was positioned close to the tree being viewed, also approximately 6m away from the camera. The device was part of a weather measurement system, equipped with a short-range radio link that allowed remote capture of the anemometer data, and the relevant datasheet is available in the Appendix.

3.5. Summary

The location was chosen to be representative of a typical scene observed by a camera in a suburban setting. It is recognised that in any location of this type, weather properties will differ from those associated with a more open location. For example, sheltering will result in reduced wind amplitude, but also in some channelling of wind and in expected turbulence.

Chapter 4 – Rain

4.1. Objective

In this section, the goal is to show effective usage of Computer Vision techniques to provide a qualitative but robust estimate of rainfall intensity.

4.2. Effects of Rain on Images

Rain gives a "streaking" effect on the processed image, as documented in chapter 2. The objective is to derive an estimate of rainfall by counting frames with "valid "streaks that correspond to raindrops within a size range and in a relevant observation region.

In an environment such as a garden where the background is largely composed of trees and other vegetation, movement of leaves and other features under windy conditions will translate to noise on the image. Rejection of this noise requires setting a threshold for the width and length of streaks, and the implicit filtering associated with these thresholds is outlined in section 4.3.

4.3. Experimental Conditions and Theoretical Analysis

In the setting used, the horizontal camera angle was 11.5 degrees. The field of view was as shown in chapter 3.

The drop size property may be considered as distributed approximately following the Marshall-Palmer graph as in figure 2.1(a). This is modelled using the drop density relationship of equation 4.1 (Uijlenhoet (2001)), with the modelling being graphed as in figure 4.1(a).

$$N_V(D) = N_0 \exp(-\Lambda D), \qquad N_0 = 0.008, \qquad \Lambda = 4.21\mathrm{R}^{-0.21}$$

Equation 4.1. *D is the drop diameter in mm, with R the rainfall rate in mm/hr*

Given considerations of noise immunity, particularly in the garden location under windy conditions, it was appropriate to concentrate on wider streaks. An arbitrary threshold of 50% fill in a track of 5 pixels width, i.e. an average width of 2.5 pixels, was established. This has filtering properties in relation to drop diameter and distance combinations, and these are analysed below.

The second filtering effect used – on grounds of noise immunity and to avoid counting partial streaks – was to require that the streak had length greater than one third of the vertical pixel count, or 160 pixels. This corresponds to 50mm in actual length per metre of distance from the camera. This only becomes the dominant constraint at drop sizes exceeding 5.5mm, and is therefore not the dominant constraint in respect of the relevant drop sizes studied.

The filtering effect associated with streak width means that drops corresponding to less than 2.5 pixels in average image width are not captured. For a 5mm drop this would equate to streaks beyond 6.4m from the camera, and in the case of a 1.5mm drop the corresponding figure is 1.92m. There is also a constraint on capture of drops near the camera, due to effects including distances below the minimum focus distance, shading effects causing light reduction, etc. The aggregation of these effects is modelled as a loss of focus below 1.8m, as shown in figure 4.1(b).

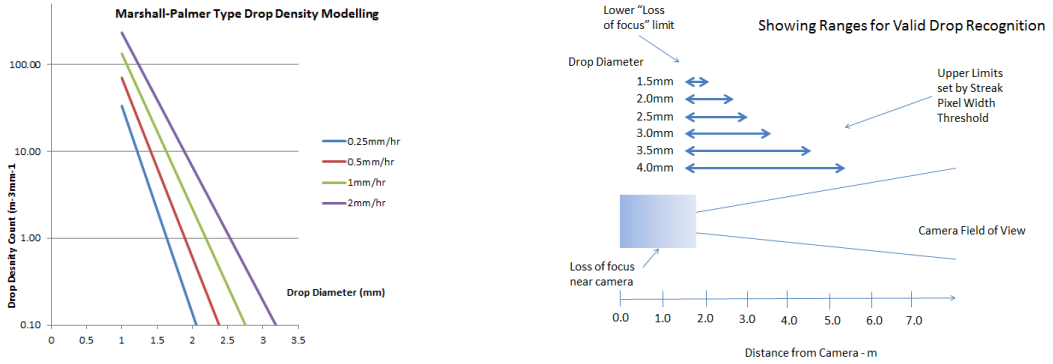

Figure 4.1(a) – left – showing curves derived from equation 4.1, and Figure 4.1(b) showing the drop capture range from the camera as a function of drop width.

The Marshall-Palmer type distribution as cited relates to a volumetric distribution, and this needs to be mapped to a distribution of drops passing a datum level – such as would correspond to a horizontal line in the centre of the camera viewing area. As a result, adjustment is needed for the terminal velocity of drops. It is also necessary to calculate the relevant area of drop capture, and adjust for this figure.

Arriving raindrops may be expected to follow a Poisson-type distribution. If the key metric is "frames per second with at least one streak" then this figure is computed as below.

Rainfall Rate (mm/hr)	0.25	0.5	1.0	2.0
Qualitative in "Irish Context"	Light	Moderate	Heavy	Very Heavy
Valid Streaks per Second	0.44	1.76	6.05	18.3
Valid Streaks per Frame	0.0147	0.0587	0.202	0.611
Expected Frames per second with at least one streak	0.44	1.71	5.48	13.7

Table 4.1 – Expected frames with streak rate as a function of rainfall level

The approximations inherent in this analysis are acknowledged, primarily relating to assumptions relating to camera minimum depth of field and relevance of the drop distribution model to local conditions. The minimum depth of field point could be addressed with precision optics, and a more open location would be preferred for elimination of other near-camera effects. A proposed topic for further work would include studies on accuracy of quantitative estimates derived in such a setting.

However, it is important to establish this theoretical understanding for correct interpretation of the experimental data.

The speed of the drop will also be influenced by windspeed. In the parallel work on analysis of wind, it was noted that average wind speeds of up to 40km/hour could be observed in the relatively-sheltered location chosen. This results in the path of raindrops deviating materially from the vertical (see for example the images as in figure 4.2(a) and 4.2(b)), and this aspect is reflected by adjusting the pixel width threshold for the relevant angle. This is implemented in the image processing code.

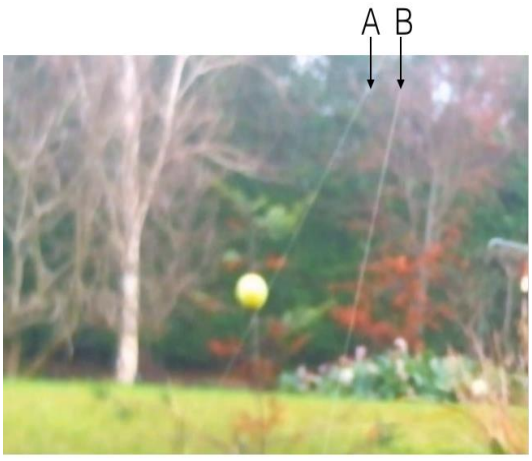

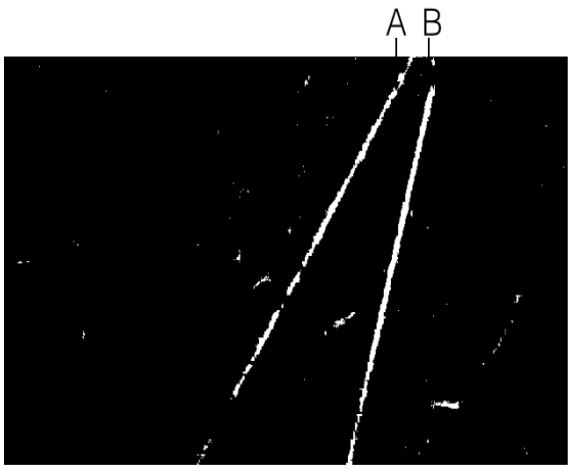

Figure 4.2(a) – left – showing wind-blown rain streaks ("A", "B") on full image as captured, and figure 4.2(b) showing image after processing (subtraction of prior image and threshold-based conversion to binary image)

4.4. Experimental Setup

This was as used for measurement of wind. A fixed-position camera was used, housed in a plastic cover to protect it from the rain.

4.5. Process

The process undertaken in analysis of a video for rain is described as follows and summarised in Chart 4.1.

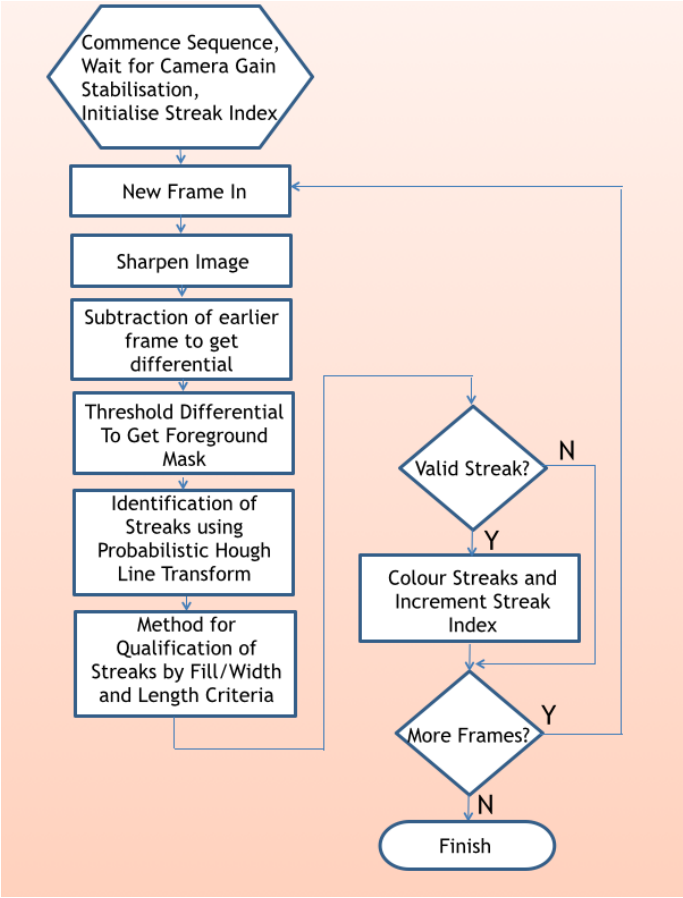

Chart 4.1. – Rain Image Analysis Process Flow Chart

A video recording was undertaken, lasting 60 seconds in duration and containing 1800 frames. This duration allows analysis of a meaningful number of frames, permitting an

accurate representation of the weather in the scene. Every frame in the video sequence was processed and analysed.

Due to the refraction of light from raindrops, a falling raindrop will produce a streak of enhanced light level as compared with the rest of the image. Conversion to the HLS representation allows capture of the luminance information. By using this luminance channel, it allowed the brighter streaks stand out further against the darker background. Figure 4.3(a) is an example of an RGB representation of a frame and Figure 4.3(b) is the corresponding luminance representation of the same frame.

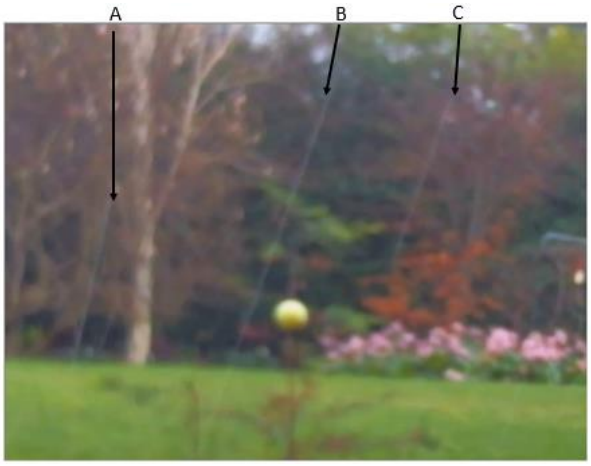

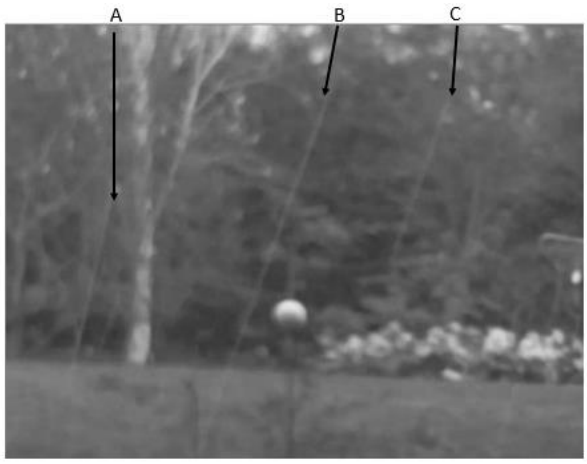

Figure 4.3(a) – left – RGB representation and Figure 4.3(b) – right – associated luminance component of an image

It was then necessary to generate a foreground mask which contained the rain streaks and any other moving objects in the scene. Background subtraction was used to create the mask. Background subtraction calculates an initial foreground mask by performing a subtraction between the current frame and a background model, in this case the previous frame in the video, before a threshold is applied to the foreground mask. By applying a threshold to the foreground mask, pixels which do not meet a desired intensity value are removed.

Figure 4.4(a) – left – showing the "previous" image and Figure 4.4(b) showing the "current" image, with a visible streak effect

An example of background subtraction of Fig. 4.4(a) and (b), with the corresponding foreground mask as shown below in Fig. 4.5(a) and (b).

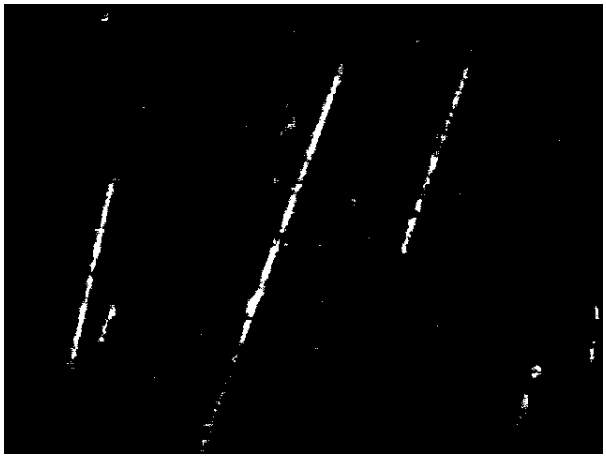

Figure 4.5(a) – left – showing the background subtraction of the two frames of Figure 4.4
Figure 4.5(b) – right – the foreground mask after thresholding the background subtraction

The foreground mask will include most moving objects in a scene provided they stand out from the background. If it is a windy day, movement of trees will be captured. If an animal or human should walk through the scene then they too will be included in the foreground mask. Therefore it was necessary to differentiate rain streaks from other objects or noise present. As noted previously, a relevant rain streak in a frame will fall in one direction and due to the exposure time of 33ms it can be expected to be at least one third the height of

the screen. Therefore, it was appropriate to look for straight lines which measured at least 160 pixels in length.

A Hough Transform technique was used to detect the presence of straight lines in the foreground mask. Hough Transform is a technique, the purpose of which is to find imperfect instances of objects within a certain class of shapes by a voting procedure (Duda, R. and Hart, P., 1972). Both the Standard Hough Line Transform and the Probabilistic Hough Line Transform can be used to identify the presence of lines. For this process, the Probabilistic Hough Line Transform is used because it identifies line segments as components of lines, with a tolerance for small gaps. By using this method, it was possible to disregard line segments less than 160 pixels in length and also segments that had gaps of greater than 20 pixels. A vector of line segment coordinates was returned by the method.

Once the line segments in the mask had been identified it was important to assess if each line segment corresponded to a meaningful rain streak and a post processing routine was used for this purpose. Each line segment was verified individually by analysing the segment in the mask that corresponded to the area covered by the line. A criterion for acceptance of a line segment as a constituent of a valid streak was 50% fill factor in a box of width 5 pixels (the centre of the line +/- 2 pixels) and of length 160 pixels. An example of an invalid line segment being rejected is shown in figure 4.6. In the case of thick streaks, a number of adjacent line segments may meet this criterion. However, as the metric is “frames with at least one streak”, this additional information is redundant.

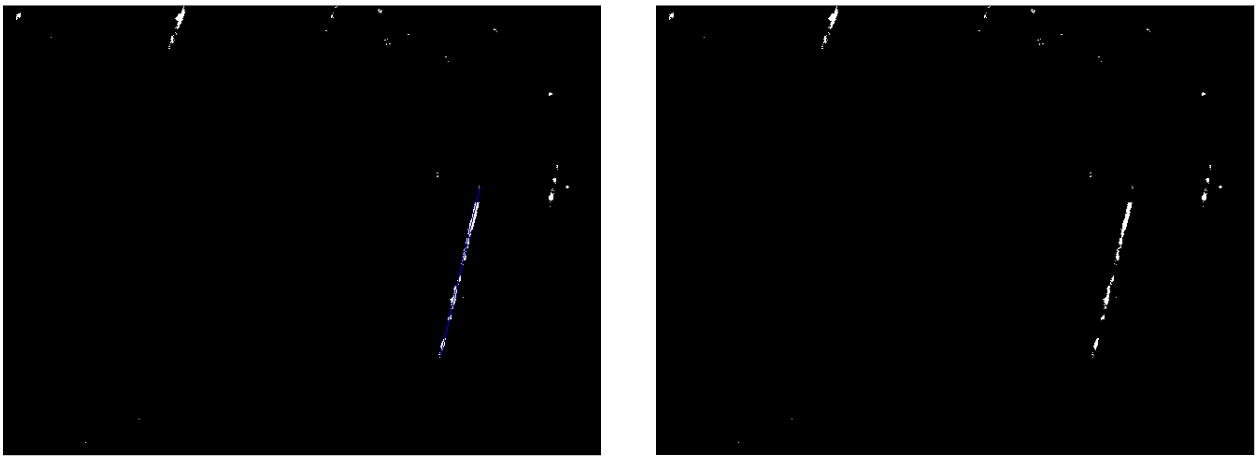

Figure 4.6(a) – left – A line segment identified by the Probabilistic Hough Line technique but removed – right - due to the fact that the area bordering by the line has a fill of less than 50 %

Once all the valid rain streaks had been identified it was appropriate to highlight them. Each valid rain streak in the foreground mask was coloured in red using the method *colourStreak*. An example of this is shown in figure 4.7.

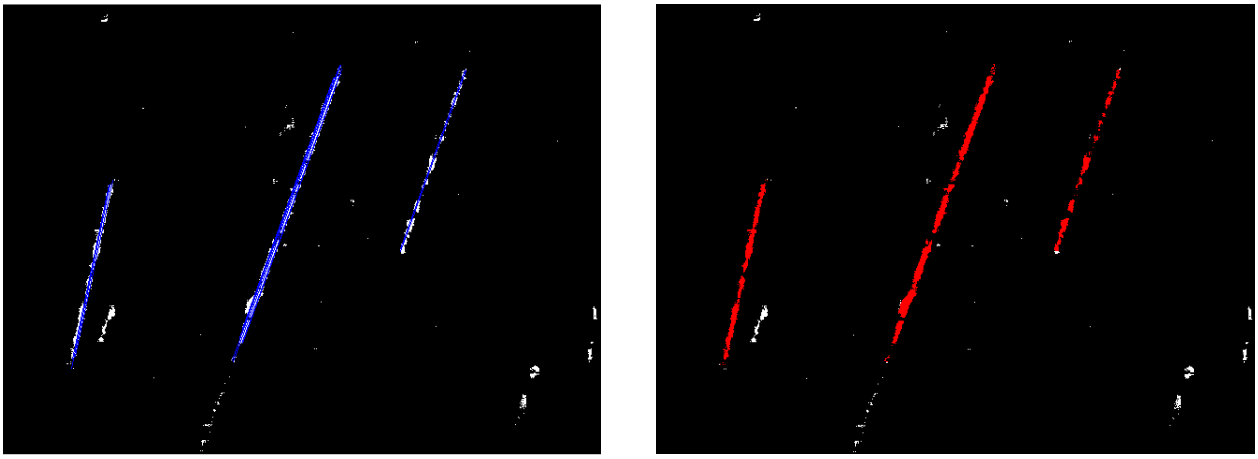

Figure 4.7(a) – left – the foreground mask from Figure 4.5(b) with the individual line segments identified coloured in blue and – right – the foreground mask with "valid" rain streaks coloured in red.

The intensity of the rain was estimated by identifying the number of frames in the video containing valid rain streaks and then calculating the percentage of these in the complete video sequence.

4.6. Results from Analyses of Video Recordings

Twenty seven video recordings were used for validation, all using consistent camera placement and settings. These were each of 60 seconds duration, thus containing 1800 frames.

The inferred rain intensity deriving from each recording was categorised into one of five bands as shown in Table 4.2. This categorisation is informed by the theoretical calculations based on the Marshall-Palmer distribution, as verified by human observation.

Percentage of Frames with Rain Streaks	Qualitative Observation
<1%	No/Minimal Rain
1-4%	Light Rain
4-10%	Moderate Rain
10-20%	Heavy Rain
20%+	Very Heavy Rain

Table 4.2. Interpretation of Rainfall data.

The data is optimally presented using the matrix below, where the predicted band from measurement appears in columns, with rows corresponding to bands derived from qualitative human observation.

		Calculated Category				
		No/Min	Light	Moderate	Heavy	Very Heavy
Observed Category	No/Min	7	0	0	0	0
	Light	0	4	1	0	0
	Moderate	0	1	5	2	0
	Heavy	0	0	0	2	2
	Very Heavy	0	0	0	0	3

Table 4.3 Results matrix for rain observations

These results are seen to be consistent with the fundamental analysis work as presented earlier in this chapter.

4.7. Analysis

4.7.1. Strengths

This work demonstrated successful usage of Computer Vision techniques for capture and processing of rain information. This information appears as “streaks” on images, with the frequency of frames with at least one valid streak presenting an estimate of rainfall intensity.

The calculations using established rainfall data and drop distribution size showed broad consistency with the measurements using Computer Vision techniques. The results were also comparable with published historical data for the region in question.

This work builds on prior work as cited in chapter 2, including treatment of off-vertical streaks as associated with wind effects and by usage of an alternative filtering approach to address the "noisy" environment associated with a garden setting.

4.7.2. Limitations

The principal limitation relates to effects close to the camera. Here, loss of camera focus, shading of incident light from raindrops and physical shading of raindrops are likely to produce the most significant errors. This could be addressed in part by usage of advanced optics giving a more precise depth of field.

4.7.3. Future Work

The vision of an "Electronic Rain Gauge" using Computer Vision techniques has been cited by Garg and Nayer (2007). Obtaining accuracy levels suitable for quantitative meteorological analysis would require more advanced optics in an open location, and verification of performance under these conditions would be of value.

4.8. Conclusions

Direct observation of falling raindrops is seen as allowing inference of rain intensity. The techniques presented provide enhanced noise immunity, and the effects of the filtering measures are quantified in terms of their effect on the range of drop sizes whose patterns are captured.

A foreground mask is derived using subtraction and thresholding approaches. Usage of the Probabilistic Hough Line Transform on the resulting binary image gives an initial indication of streaks, with this information being further processed to validate such streaks and thus infer rainfall intensity.

Chapter 5 – Light

5.1. Objective

Obtaining a qualitative metric as to weather conditions being “overcast” or “bright” is the goal of this chapter.

5.2. Effect of Light on Images

Light effects can be observed primarily through the following:

Shadow Observation - This can be feasible when the field of view of the camera can identify clear shadow patterns at all relevant times of the day, with a sundial-type element being an example of an object that could generate a consistent shadow pattern.

Overall Light Level - The field of view was such that consistent analysis of shadows was not possible, and thus the focus was very much on deriving observations based on overall light level within the image.

Additional techniques could involve analysis of incident sunlight on high branches or making sky colour evaluations, but these approaches are challenged by the need to get robust results at various times of day and under different conditions of wind.

5.3. Range of Light Intensity

The image as recorded by the camera corresponds to the signal on the photodetector elements, as adjusted by an internal automatic gain control. This gain control is essential if the dynamic range of the camera is to be adequate for the wide range of typical user requirements. For example, a dynamic range requirement of approximately 10^6 is expected if the camera is to be able to resolve images under low light conditions whilst also avoiding saturation under very bright conditions, such as for example those associated with direct sunlight impinging on the image.

This need for significant dynamic range can be seen from Figure 5.1.

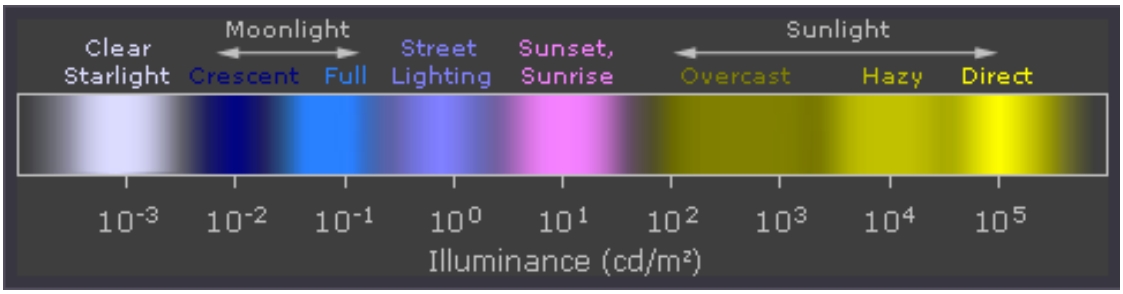

Fig. 5.1. Approximate values for commonly encountered light sources (Cambridge In Colour)

For the purpose of measurement of light levels, selecting a suitable fixed gain value, thus disabling the automatic gain control functionality, is necessary. This can be undertaken under software control using the Logitech camera API.

This study is associated with daylight conditions, and therefore values materially below approximately 100cd/m^2 are of limited relevance. Given the shading of the much of the image, making the distinction between light levels corresponding to direct and hazy sunlight (i.e. levels above approximately 3000cd/m^2) is less relevant. The gain control setting is thus selected so that the camera captures this range of light levels, allowing a margin for saturation effects.

It is then possible to categorise light levels into qualitative bands approximately defined as follows:

Very Overcast	<600cd/m2
Overcast, Heavily Shaded	600cd/m2 to 3000cd/m2
Bright	>3000cd/m2

Table 5.1 – Light level categorisation

This type of “banding” of light levels also correlates with material from Schlyter (2004) summarised as in figure 5.2.

	Stellar magnitude	Illuminance Lux
Sun overhead	-26.7	130000
Full daylight (not direct sun)	-24 to -25	10000-25000
Overcast day	-21	1000
Very dark overcast day	-19	100
Twilight	-16	10
Deep twilight	-14	1
1 Candela at 1 meter distance	-13.9	1.00
Full Moon overhead	-12.5	0.267
Total starlight + airglow	-6	2E-3
Total starlight only	-5	2E-4
Venus at brightest	-4.3	1.4E-4
Total starlight at overcast night	-4	1E-4
Sirius	-1.4	1E-5
0th-mag star	0	2.7E-6
1st-mag star	+1	1.0E-6
6th-mag star	+6	1.0E-8

Fig. 5.2. Light levels associated with natural lighting conditions (Schlyter, 2004)

The SI unit of light intensity is candela/m^2, as used earlier, and this is synonymous with the "Lux" term which is used more conventionally.

5.4. Light Effects

With the gain control set at a fixed level, the camera may operate as a light-meter, by using an approach such as getting the average luminance value of the pixels within the image. The relationship between this value and the light level needs to be determined. This is best established using calibration against a commercial light-meter.

5.5. Approach

The experimental setup involved usage of the same Logitech C615 HD Camera as in previous work. The automatic gain control within the camera was disabled using software control. A fixed gain setting was chosen, that visibly corresponded to saturation onset under strong light levels associated with direct sunlight and this gave a readable image under twilight conditions (essentially at the time of documented sunset).

This was calibrated against a commercially available light-meter, as documented in Chapter 3. The reading is in the SI unit of lux, synonymous with cd/m^2.

The light-meter is equipped with a hemispherical-type light integrator with a wide field of view. In order to optimise the calibration, the field of observation of the light-meter was constrained to make this approximately similar to the FOV of the camera. This involved placing the light-meter at the back of a section of tubing, which could then point in the same direction as the camera.

The tubing reduced incident light on the light-meter, and it was necessary to correct for the reduction in viewing angle associated with this tubing. Initial studies gave a correction factor requirement of 4:1 for when the tubing was place on the light-meter.

The calibration process initially involved deriving a brightness index for the image. A simple index can be derived by using the average luminance value from the pixels in the image. 120 images taken under various light levels were characterised in this fashion, and the index values plotted versus the adjusted lux figures from the light meter.

This gave a logarithmic plot as in figure 5.3, with a correlation (R^2) figure of .94. The logarithmic aspect ties in well with expectations, as many human senses tend to be correlated logarithmically to physical phenomena.

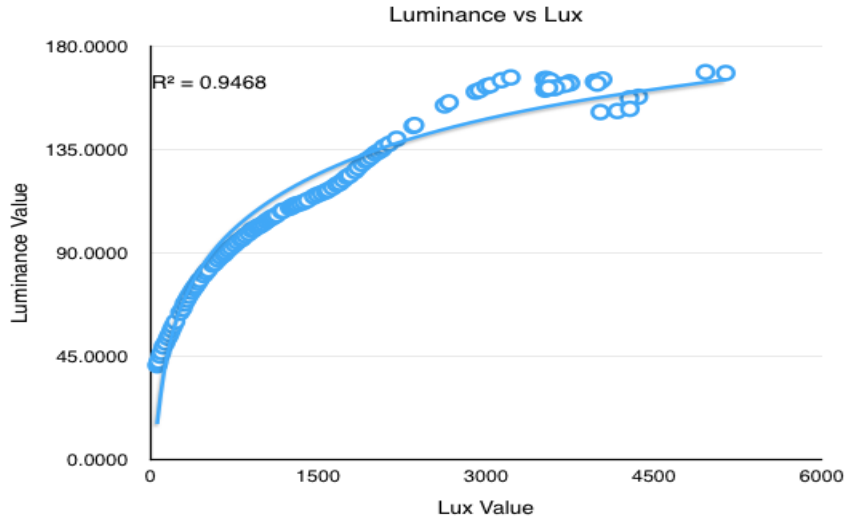

Figure 5.3. A correlation curve with a logarithmic trendline representing a Luminance index vs. measured Lux values

5.6. Results

Each of the images used for camera calibration is also checked visually as corresponding to the relevant band as implied by the values, with a number of representative images as shown below.

Image #	Index (average pixel luminance value)	Implied Lux Value	Banding	Human Observation
1	163	4874	Bright	Bright

2	131	1869	Overcast/Heavily Shaded	Overcast/Heavily Shaded
3	91	580	Very Overcast	Very Overcast
4	151	3421	Bright	Bright
5	130	1845	Overcast/Heavily Shaded	Overcast/Heavily Shaded
6	99	730	Overcast/Heavily Shaded	Overcast/Heavily Shaded

Table 5.2 – Sample result categorisation

1.

2.

3.

4.

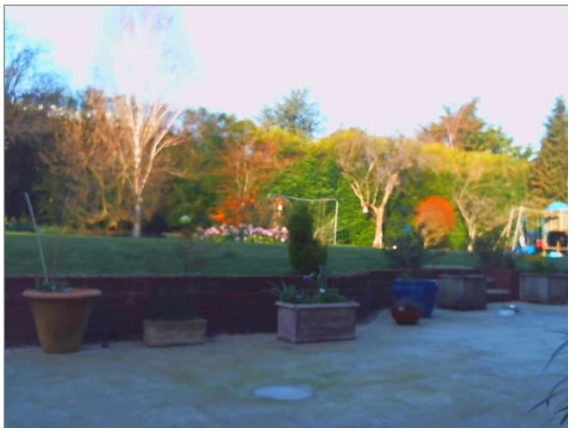

5.

6.

Figure 5.4 – A number of images with varying levels of light intensity

5.7. Analysis

Obtaining information as to light conditions is a further capability, provided in addition to the key areas of wind and rain analysis.

5.7.1. Success

Sufficient correlation between an average luminance value and readings on a commercial light-meter was obtained, with the camera shown as capable of operating relatively consistently as a light-meter over the range of interest.

5.7.2. Limitations

The desire was to maintain the camera position similar to that used in other sections of this work but adjust the zoom to get a full view of the garden, and this (see map from earlier) corresponded to the camera looking west. The garden setting was largely shaded by buildings up to mid-morning, and some shading was provided by trees later in the day. A consequence of this was reliance on overall light level analysis rather than being able to use shadow type techniques.

Additional information as to presence of direct sunlight could be obtained by ensuring that sundial type objects giving a reliable shadow pattern were present within the field of view, but this is also dependent on shading effects at various times of day.

5.7.3. Future Work

An elevated camera position would give greater shadow observation opportunity but again this is subject to limitations associated with the location chosen. Sky colour analyses may also be feasible in determining the level of cloud cover, which is naturally associated with determination of overcast or bright conditions.

5.8. Conclusions

In this chapter, it is seen that the camera can be readily used as a light meter and that this can provide valuable complimentary material to wind and rain information.

Chapter 6 – Wind

6.1. Objectives

This chapter is concerned with studying movement of those objects which can be correlated with wind characteristics. Experimental work establishes the extent of such correlation and evaluates two Computer Vision techniques, discussing also their limitations.

"Wind" is the movement of atmospheric air. This movement is in itself naturally invisible. Therefore it is necessary to infer wind characteristics from looking at the effect of wind on objects within the field of view.

6.2. Effect of Wind on Images

Initial observation of differences brought about by increasing wind levels suggests evaluation of effects such as:

- Flag-type or Windsock-type behaviour where an object is at rest in low wind conditions and where the elevation angle of the object is broadly correlated with wind strength over a realistic range. For example, commercial windsocks for aviation are specified so as to give an initial indication of wind direction from speeds of 5.6km/hr to being fully extended at windspeeds of 28km/hr (FAA, 2004).
- "Tilt" type effects, such as when a small tree can be observed as bending in wind.
- "Fluttering" and comparable examples of harmonic motion in vegetation, typically movement of leaves and smaller branches.

There can also be transient effects such as those observed in chapter 4, where wind speed could be inferred from the angle of falling raindrops. This work has had value in giving insight into the turbulent nature of wind within the field of view, but is of course unsuited to measuring wind in the absence of rainfall and is therefore not considered further.

In this chapter the tilt and leaf-movement approaches are considered and appropriate techniques are analysed, with a commercial anemometer used to provide calibration.

The experimental setup was outlined in chapter 3. The location as shown in figure 3.1(b) is sheltered, which indicates that windspeeds may be materially lower than for example values published by meteorologists based on measurements at stations in open terrain. Local topographical features will also cause aspects of turbulence and channelling of wind.

6.3. Wind Characteristics

6.3.1. Initial Analysis

Within the field of view normally available to a camera in a "conventional" setting, it is unlikely that a windsock or flag would be present. One thus needs to focus on tilt effects and on the fluttering and other harmonic motion effects associated with the impact of wind on vegetation.

As previously mentioned, the impact of wind on tree movement has been studied in some detail in James et al (2006) and Mayer and Schindler (2009). These works present comprehensive analysis of tree mechanics, considering tilt of the overall tree and the swaying of branches with their individual harmonic components. These harmonic components naturally add together at the peripheries of the tree, meaning that leaf movement is likely to be less well defined than movement of a more central part of the tree.

The spread in tree characteristics, and the variation in loading due to the presence or otherwise of leaves, confirms the difficulty in deriving a theoretical basis for linking observed tree movement to wind loading. The principal observation deriving from this review is that there may be value in seeking to identify absolute movement of two components. Movement of the larger trunk-type elements may be described as "Central Movement" and the quasi-random movement associated with elements at the peripheries can be described as "Peripheral Movement".

6.3.2. Central Movement

The "tilt" movement at the centre of the tree is naturally more limited, and the trunk may be less visibly differentiated from the background. As a result it was appropriate to identify a small tree whose movement would be more evident, and its location several metres from the camera allowed movement to correspond to a meaningful number in terms of pixels

within the image. The image could be enhanced by affixing a small "target" to the tree, a yellow tennis ball, and this could allow tracking using established Hough Circle techniques.

6.3.3. Peripheral Movement

This work used an Optical Flow technique developed by Gunnar Farneback to study the movement of features including branch extremities. This complements the measurement of the tilt (i.e. displacement of the target at the centre of the tree).

6.4. Approach

Chart 6.1 represents the process used to capture and analyse wind information.

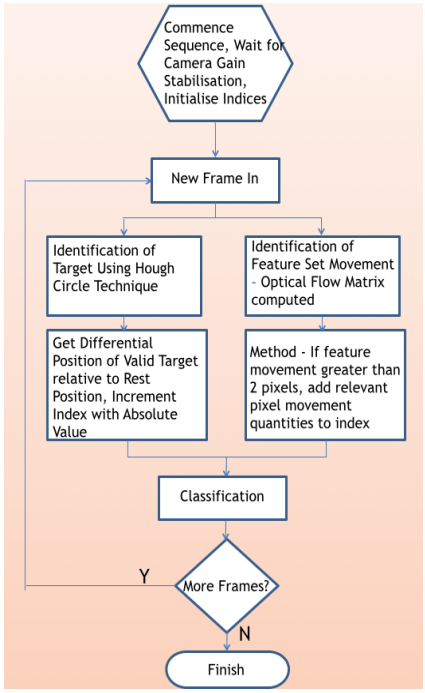

Chart 6.1. A flowchart depicting the overall analysis of wind

Each video sequence lasted approximately 30 seconds. The videos were of this length so as to be of similar duration to the observed data capture interval of the anemometer. It is important to allow the camera's automatic gain control become operational and thus the first 50 frames of the video sequence is skipped. Every 8th frame in the video was analysed, giving a sampling period of 270ms, which was considered appropriate given the visibly slow movement of the tree elements.

The initial technique used was the Optical Flow algorithm of Gunnar Farneback (2003). Optical Flow is the pattern of apparent motion of image objects between two frames caused by the movement of object or camera. The Farneback function compares a current frame versus a previous frame, estimates a displacement between the same features in each frame and finds the "flow" for each pixel of the previous image. The Farneback algorithm calculates the Optical Flow such that -

$$\mathtt{prev}(y,x) \sim \mathtt{next}(y + \mathtt{flow}(y,x)[1], x + \mathtt{flow}(y,x)[0])$$

Equation 6.1 Farneback algorithm (OpenCV)

As outlined in the experimental setup, the camera is programmed to focus upon a small tree in the centre of the garden that is likely to be affected by wind. The Farneback method is used to calculate the Optical Flow of the visible features of the tree, returning a matrix of feature movement data. Only the area of interest, the area immediately around the tree, is analysed, thus avoiding movement of extraneous objects.

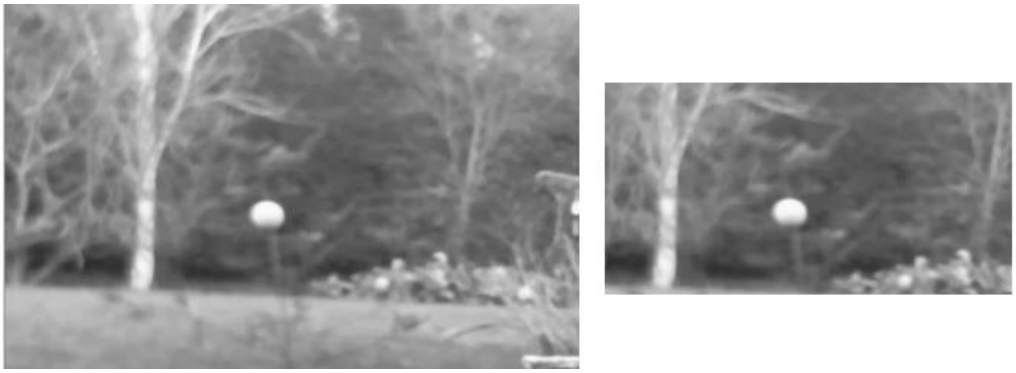

Figure 6.2. A screenshot of the field of view for the camera (left) and the area that is analysed for Optical Flow (right).

The method *measureOpticalFlow* was then used to process this data and derive an overall index representing the movement of the relevant pixels. If a feature has moved by more than two pixels (corresponding to approximately 4mm at the tree) along the x-axis then it is recorded as movement. The movement of all the pixels in the area of interest is aggregated. A portrayal of Optical Flow analysis is shown in the figure below.

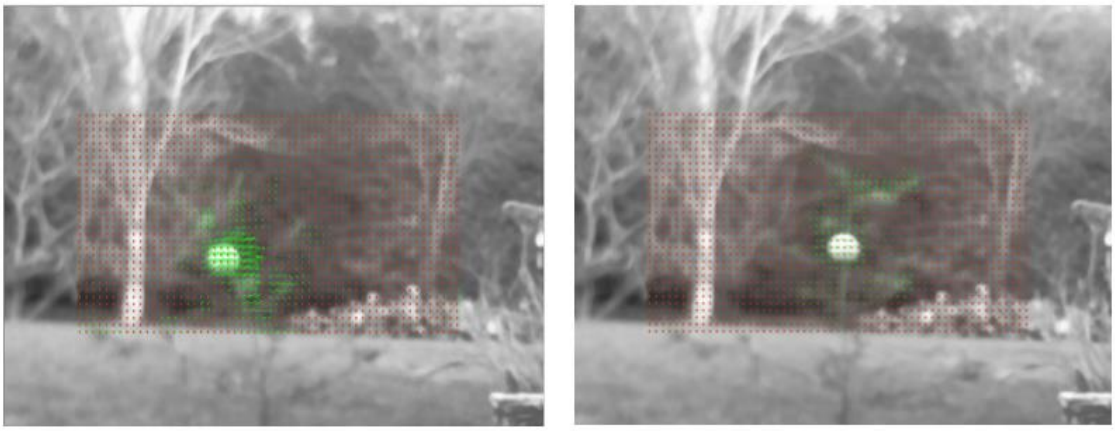

Figure 6.3. Examples of frames with Optical Flow shown. In this case the green arrows indicate movement.

The Optical Flow technique results can be correlated with wind speed and this is undertaken subsequently in this chapter. The Optical Flow technique may also be of value in deriving an index related to wind variability. This may be computed by processing the data using moving average filtering as illustrated in the Results section.

The second technique used is the Hough Transform to identify a target which is present on the tree. As outlined in Chapter 4, Hough Transform is a technique whose purpose is to find imperfect instances of objects within a certain class of shapes by a voting procedure. Hough Circles was used because the target in question was a tennis ball. Figure 6.4 shows examples of frames with the target highlighted by Hough Circles.

Figure 6.4. Examples of frames where the target has been identified by Hough Circles

By using the location of the tennis ball in the image, it was possible to identify the displacement of the tree caused by wind. In order to calculate this displacement, the target position under calm conditions was established, and this was at the 315 pixel along the x-axis. Therefore it was possible to calculate the displacement caused by wind as the absolute value of 315 minus the x-value of the target in a frame.

Fig. 6.5. One frame where the target is at rest (left) and another where the target has moved from the rest position (right). The displacement caused by the wind in the right frame is of 88 pixels.

The average absolute displacement of the target is computed, rejecting any occasions where the circle is not identified correctly. It is then necessary to correlate this index information and the Optical Flow information against the wind speed as estimated by the commercial anemometer.

This information is plotted in Figure 6.7(a) and Figure 6.7(b), comparing both metrics with the anemometer reading.

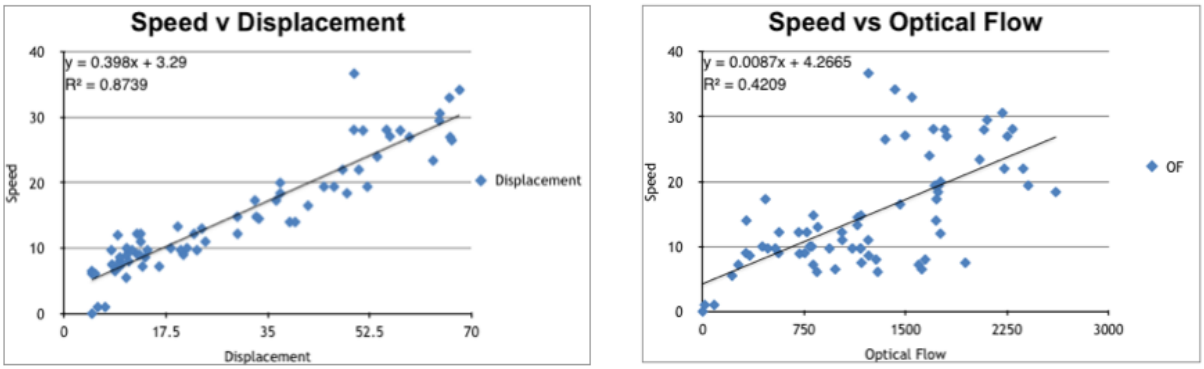

Figure 6.7(a) Scatter chart indicating the speed vs. mean target displacement (left) and Figure 6.7(b) Scatter chart indicating the speed vs. Optical Flow index (right)

The correlation value of 0.87, between the displacement metric and the anemometer speed indication, is seen as realistic given the conditions associated with observation of wind effects. Specifically, a tree is not particularly well quantified in terms of its dynamic behaviour. Some outliers with respect to the trend-line are seen but these are unlikely materially to affect qualitative results presentation.

There was a less pronounced correlation with the broader movement characterisation as captured using the Optical Flow technique. This gave a result as shown with a correlation coefficient (R^2) of approximately 0.42.

The trendline of Figure 6.7(a) gives the relationship of equation 6.2:

$$\text{implied_speed} = (0.398*\text{totaldisplacement}/(\text{framecount}-\text{nocircle})+3.29)$$

Equation 6.2. Equation to calculate the estimated speed of the wind in the video.
totaldisplacement is the sum of all of the displacement values from frames analysed.
framecount is the number of frames analysed.
nocircle is the number of frames which no target was identified.

The process for estimating wind speed in a video sequence involves deriving this mean absolute displacement and inferring a wind speed based on the equation above. The wind level is then associated with one of a number of bands such as shown below (Table 6.1).

Implied Bin	Speed – km/hr
No/ Minimal Wind	<4
Light Wind	4-6
Light/Moderate Wind	6-9

Moderate Wind	9-13
Moderate/Strong Wind	13-18
Strong Wind	18-23
Strong/Very Strong Wind	23-28
Very Strong Wind	28+
Table 6.1 Wind Bands and associated speeds	

This approach is influenced by the Beaufort scale (Met Eireann, 2014). The absolute values of the Beaufort scale apply to exposed locations and thus are not directly applicable to a sheltered location such as considered here.

The estimated wind speed, wind type and wind variability are displayed in the top left corner of the image in order to deliver real-time information after analysis of 30 frames.

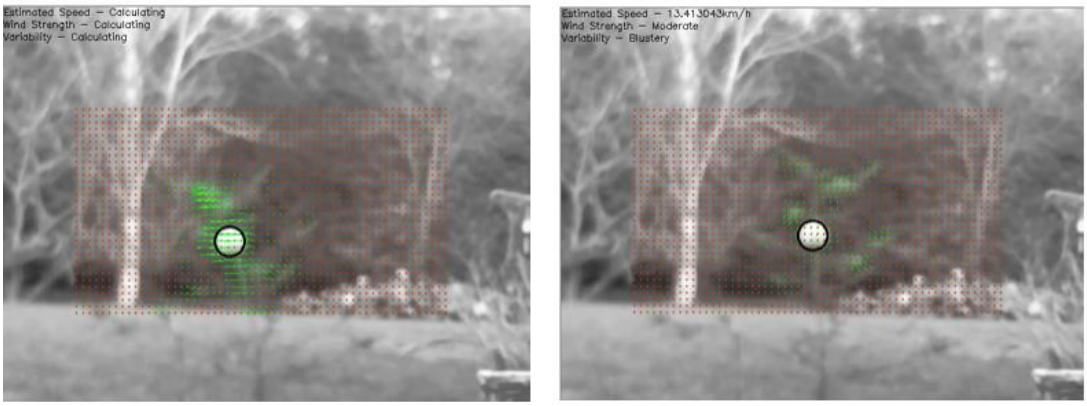

Fig. 6.6 Two frames from the program, the left frame still calculating while the right frame showing real-time information

6.5. Results

Results from 70 videos have been collated in the material provided below, with the key metric being absolute displacement – as obtained using the Hough Circle technique

The average windspeed data is best presented qualitatively using the data binning approach, with examples as follows:

Video Number	Estimated Speed km/h	Estimated Band	Anemometer Speed km/h	Anemometer Band	Band Consistency Within +/- One Band	Correct Band
4	11.7	Moderate	10	Moderate	Yes	Yes
9	16.6	Moderate/Strong	14.5	Moderate/Strong	Yes	Yes

17	8.5	Light/Moderate	11	Moderate	Yes	No
20	7.0	Light/Moderate	7.2	Light/Moderate	Yes	Yes
38	16.4	Moderate/Strong	14.8	Moderate Strong	Yes	Yes
51	6.1	Light/Moderate	1	No/Min	No	No
63	29.8	Very Strong	26.5	Strong/Very Strong	Yes	No

Table 6.2 Data presented for a number of videos analysed

The data was analysed and 69 of the 70 samples were found to exhibit a similar band ± one band, when comparing the estimated speed versus anemometer speed, and with 50% giving an identical band reading (see Table 6.3).

		Predicted Category							
		No/ Minimal	Light	Light/ Moderate	Moderate	Moderate /Strong	Strong	Strong/ V. Strong	Very Strong
Anemo-meter Derived Category	No/Minimal	0	2	1	0	0	0	0	0
	Light Wind	0	0	1	0	0	0	0	0
	Light/ Moderate	0	3	10	1	0	0	0	0
	Moderate	0	0	11	7	1	0	0	0
	Moderate/ Strong	0	0	0	2	5	3	0	0
	Strong	0	0	0	0	0	6	2	0
	Strong/Very Strong	0	0	0	0	0	0	3	3
	Very Strong	0	0	0	0	0	0	5	4

Table 6.3 Results Matrix for Wind Analyses

It is recognised that usage of the “target” approach is likely to show increased sensitivity to displacement across the field of view, rather than in-line with the camera axis. This represents an inherent limitation on the accuracy of this approach.

6.6. Analysis

6.6.1. Successes

This work represents a valuable approach to inferring wind characteristics from movement within the field of view. In the absence of windsock or flag elements, tree movement

properties were analysed and offered an acceptable proxy for wind speed estimation. To get consistent results, it was necessary to include the tilt of the tree. This provided correlation to a level consistent with qualitative estimation of wind speed in meaningful bands.

6.6.2. Limitations

The principal limitations will be associated with local topographical features, namely sheltering by buildings and trees, as well as ground proximity effects. These all give rise to turbulent air conditions. The measurement approach is also sensitive to wind direction, with movement across the field of view being fully taken into account whereas this is less true in the case of movement along the axis of the camera orthogonal to the focal plane. If deciduous trees are selected, then variability in leaf cover will affect results. Under instances of very strong wind, the movement of the tree can lead to a blurring effect on the ball image which may prevent recognition using Hough Circles.

6.6.3. Future Work

An implied constraint related to usage of one camera. Usage of more than one camera, with an orthogonal axis, could allow assessment of wind direction. Benchmarking this work in more exposed areas, ideally located adjacent to a weather station, would also allow better calibration. Usage of a purpose designed instrument, consisting for example of a mass on a flexible pole with appropriate damping, could potentially allow usage of Computer Vision techniques for full quantitative analysis.

6.7. Conclusions

The work associated with this chapter points out opportunities and limitations associated with using Computer Vision techniques for wind analysis. Unlike the treatment of rain where the trails associated with individual drops can be characterised, this subject area relies on movement of vegetation. However, this movement of vegetation under the influence of wind is an area that has historically received limited research attention.

Wind effects are categorised into those which primarily are associated with “tilt” or displacement of a central tree section, and others that primarily affect the periphery. The tilt effects showed correlation results with windspeed that were considered acceptable in the context of provision of qualitative information. The peripheral effects are significantly

more complex in terms of analysis, given the multiplicity of mechanical resonances in branches. An indicator derived from movement of peripheral elements showed low correlation with windspeed but could be of value in giving a metric related to variability.

Chapter 7 – Conclusions

7.1. Overview

The studies as outlined in chapters 4-6, along with conclusions presented in each chapter, have confirmed that Computer Vision techniques can be of value in providing at least qualitative information as to weather characteristics in an automated fashion.

In the case of rain analysis, there is an expectation that useful real-time quantitative data could be obtained if more complex optics could be included in cameras, with some limitations associated with wind-blown rain. A good correlation was however observed between metrics deriving from filtered streak analysis with human perception of rain levels.

The correlation results obtained in the case of light show that the camera can serve as an overall light-meter, allowing determination of light levels and their interpretation along a scale from very overcast through to bright. These results could be correlated with time-of-day and time-of-year and latitude/longitude information for more precise interpretation.

The wind results showed the desirability of a "target" being present, in order to capture the "tilt" aspect in a flexible element. The mapping of incident wind direction to turbulence and displacement needs further characterisation. Otherwise, metrics associated with fluttering of leaves and other small elements were unlikely to be particularly well correlated with average or steady windspeed. The "tilt" is expected to be correlated with average windspeed, and the results indicated usable correlation levels for this approach given the range of studies undertaken. The opportunity of designing an optimised target object with defined dynamic behaviour in response to wind has also been identified.

Whilst the initial remit was concerned with daylight observations, it is reasonable to expect that artificial illumination could allow observations of raindrops and also of a wind-relevant target. This would be particularly relevant to deployment of such techniques in automated indicator systems for traffic guidance.

This work has indicated areas where CV techniques can give consistent metrics associated with weather phenomena, also setting out limitations on their usage and indicating areas relevant to future work.

Appendix 1. Light-Meter Data

Specifications:	
measuring range:	200, 2000, 20000lux (20000lux range reading x10) and 50000lux (50000lux range reading x100)
accuracy:	< 10000lux: ± 5% of rdg + 10 digits
	> 10000lux: ± 10% of rdg + 10 digits
repeatability:	± 2%
temperature	
characteristics:	± 0.1%/°C
photo detector:	one silicon photo diode with filter
power supply:	1 x 12V A23 battery (included)
measurement rate:	1.5 times per second, nominal
storage temperature:	-10°C to +60°C
photo detector	
dimensions:	115 x 60 x 27mm
dimensions:	188 x 64.5 x 24.5mm
weight:	160g

Specifications for Maplin N76CC Light-Meter - http://www.maplin.co.uk/p/digital-light-meter-n76cc

Appendix 2. Anemometer Data

Wind Speed	0 - 100mph (0 - 160kph)
Wind Speed Measurement	mph or kph (switchable)
Wind Speed Mode	average or gust (switchable)

Summary specifications for Maplin N96FY Weather Forecaster

Appendix 3. Detailed Analysis of Rainfall Characteristics

The spreadsheet as below is used for computation of rainfall characteristics. This is also provided in the attached DVD.

Rainfall Computations using Marshall-Palmer Distribution

Near Focus Limit Distance (m: 1.8
Near focal limit area: 0.324

Frames per Second: 30

Taking Marshall Palmer distribution of drops as Na

Drop Size	Distance for 2.5 pixels Resolution	Terminal Velocity	Distance in Frame	160 Pixels Limit	Area	Rate	0.25	0.5	1	2	4	mm/hr	0.25	0.5	1	2	4	0.25	0.5	1	2	4
							Marshall-Palmer Density Computation						Particles by size increment					Particles Through Horizontal Planes By Size Band				
mm	m	m/s	m				Particles per m^3 per mm width of distribution						Particles per m^3 segmented by size band (0.1mm between 1mm/hr and 2mm/hr)					Segmented by size in 0.5mm/0.1mm bands				
1.0	1.28	4.47	0.15	0.06	0.00		33.17	69.74	132.58	231.04	373.44		3.32	6.97	13.26	23.10	37.34	0.00	0.00	0.00	0.00	0.00
1.1	1.41	4.69	0.16	0.07	0.00		19.17	43.40	87.99	162.09	274.87		1.92	4.34	8.80	16.21	27.49	0.00	0.00	0.00	0.00	0.00
1.2	1.54	4.90	0.16	0.08	0.00		11.07	27.01	58.39	113.71	202.32		1.11	2.70	5.84	11.37	20.23	0.00	0.00	0.00	0.00	0.00
1.3	1.66	5.10	0.17	0.08	0.00		6.40	16.81	38.75	79.78	148.92		0.64	1.68	3.88	7.98	14.89	0.00	0.00	0.00	0.00	0.00
1.4	1.79	5.29	0.18	0.09	0.00		3.70	10.46	25.72	55.97	109.62		0.37	1.05	2.57	5.60	10.96	0.00	0.00	0.00	0.00	0.00
			0.00																			
1.5	1.92	5.48	0.18	0.10	0.04		2.14	6.51	17.07	39.26	80.68		0.21	0.65	1.71	3.93	8.07	0.05	0.16	0.42	0.96	1.97
1.6	2.05	5.66	0.19	0.10	0.10		1.23	4.05	11.33	27.55	59.39		0.12	0.41	1.13	2.75	5.94	0.07	0.22	0.61	1.49	3.21
1.7	2.18	5.83	0.19	0.11	0.15		0.71	2.52	7.52	19.32	43.71		0.07	0.25	0.75	1.93	4.37	0.06	0.22	0.66	1.68	3.81
1.8	2.30	6.00	0.20	0.12	0.21		0.41	1.57	4.99	13.56	32.18		0.04	0.16	0.50	1.36	3.22	0.05	0.19	0.62	1.68	3.99
1.9	2.43	6.16	0.21	0.12	0.27		0.24	0.98	3.31	9.51	23.68		0.02	0.10	0.33	0.95	2.37	0.04	0.16	0.55	1.57	3.90
2.0	2.56	6.32	0.21	0.13	0.33		0.14	0.61	2.20	6.67	17.43		0.07	0.30	1.10	3.34	8.72	0.14	0.64	2.30	6.99	18.27
2.5	3.20	7.07	0.24	0.16	0.70		0.01	0.06	0.28	1.13	3.77		0.00	0.03	0.14	0.57	1.88	0.02	0.14	0.70	2.81	9.32
3.0	3.84	7.75	0.26	0.19	1.15		0.00	0.01	0.04	0.19	0.81		0.00	0.00	0.02	0.10	0.41	0.00	0.02	0.16	0.86	3.63
3.5	4.48	8.37	0.28	0.22	1.68		0.00	0.00	0.00	0.03	0.18		0.00	0.00	0.00	0.02	0.09	0.00	0.00	0.03	0.23	1.24
4.0	5.12	8.94	0.30	0.26	2.30		0.00	0.00	0.00	0.01	0.04		0.00	0.00	0.00	0.00	0.02	0.00	0.00	0.01	0.06	0.39
4.5	5.76	9.49	0.32	0.29	2.99		0.00	0.00	0.00	0.00	0.01		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.12
5.0	6.40	10.00	0.33	0.32	3.77		0.00	0.00	0.00	0.00	0.00		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.03
5.5	7.04	10.49	0.35	0.35	4.63		0.00	0.00	0.00	0.00	0.00		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
6.0	7.68	10.95	0.37	0.38	5.57		0.00	0.00	0.00	0.00	0.00		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Excerpts from the spreadsheet used to estimate drop size distribution and incident rainfall characteristics

This spreadsheet is used for computation of the expected drop density as a function of rainfall level. Drop sizes are separated into bands, of 0.1mm (highlighted) between 1mm and 2mm diameter and then 0.5mm above 2mm diameter. Within each band the volumetric drop density is computed. The flow rate is then computed using this information as well as the terminal velocity of the drops, using the equation. The relevant area, considering the depth of field, the field of view and the maximum distance associated with resolution of particular drop sizes, is then used to estimate the drop rate associated with the relevant band. The expected number of drops per second is then computed by adding results for each band.

Raindrop arrival can be expected to follow a Poisson type distribution. As the most suitable metric is “frames with streaks” one needs to establish 1-P(0), where P(0) is the probability of a frame having 0 streaks. A spreadsheet as below is used for this purpose.

Frame Rate	30				
Valid Streaks per Second			0.4400	1.7600	6.0500	18.3400
Valid Streaks per Frame			0.0147	0.0587	0.2017	0.6113
Poisson Computation for (X) streaks		P(0)	0.9854	0.9430	0.8174	0.5426
		P(1)	0.0145	0.0578	0.1987	0.6024
		P(2)	0.0001	0.0016	0.0166	0.1014
		P(3)	0.0000	0.0000	0.0011	0.0207
Non-Zero Streak Count		1-P(0)	0.0146	0.0570	0.1826	0.4574
Number of Frames per Second with Streaks			0.44	1.71	5.48	13.72

Poisson Distribution for Raindrop Arrival

References

Brewer, N. and Liu, N. (2008) “Using the Shape Characteristics of Rain to Identify and Remove Rain from Video”, *Structural, Syntactic and Statistical Pattern Recognition*, Volume 5342, 2008, pp. 451-458

Cambridge In Colour, “Dynamic Range In Digital Photography”, available at - http://www.cambridgeincolour.com/tutorials/dynamic-range.htm

Cattoen, M. et al. (2005), “Compact System for Photometric Characterisation of Automotive Headlamps”, *Instrumentation and Measurement Technology Conference*, 2005. 16-19 May, 2005

Duda, R. and Hart, P. (1972), “Use of the Hough Transformation to Detect Lines and Curves in Pictures”, Communications of the ACM, Vol. 15, pp. 11-15 (January, 1972)

Farneback, G. (2003), “Two-Frame Motion Estimation Based on Polynomial Expansion”, *Computer Vision Laboratory*, Linkoping University

Federal Aviation Authority (2004), “FAA Specification for Wind Cone Assembly”, available at - http://www.faa.gov/documentLibrary/media/advisory_circular/150-5345-27D/150_5345_27d.pdf

Garg, K. and Nayar, S.K. (2007), “Vision and Rain”, *International Journal on Computer Vision*, pp. 1 – 25, February 2007

James, K., Haritos, N. and Ades, P. (2006), “Mechanical Stability of Trees Under Dynamic Loads”, *American Journal of Botany*, 93(10) 2006, pp. 1522-1530

Logitech, (2014), “Logitech HD Webcam C615”, available at - http://www.logitech.com/en-roeu/product/hd-webcam-c615?crid=34

Lu, L.C. and Ohya J. (2004), *IEEE International Conference on Multimedia and Expo (ICME)*, 27-30 June, 2004

Marshall, J.S. and Palmer, W. (1948) “The Distribution of Raindrops with Size”, *Journal of Meteorology*, Volume 5, Issue 4 (August 1948), pp. 165 - 166

Mayer, H. and Schindler, D. (2009) “Proceedings of the 2nd International Conference on Wind Effects on Trees”, Albert-Ludwigs-University of Freiburg, Germany, 13-16 October 2009

Met Eireann (2014) “Wind”, available at - http://www.met.ie/climate-ireland/wind.asp

Met Eireann (2014) “Rainfall”, available at - http://www.met.ie/climate/rainfall.asp

Met Eireann (2014), “Beaufort Scale of Wind Force”, available at - http://www.met.ie/climate/wind.asp

OpenCV - “Motion Analysis and Object Tracking” - available at http://docs.opencv.org/modules/video/doc/motion_analysis_and_object_tracking.html#id3

QuickCam Team (2007) “Pan/Tilt/Zoom Control”, Logitech, available at - http://web.archive.org/web/20101229215224/http://www.quickcamteam.net/documentation

Schlyter, P. (2004), “How Bright are Natural Light Sources?”, available at - https://web.archive.org/web/20131207065000/http://stjarnhimlen.se/comp/radfaq.html#10

Uijlenhoet, R. (2001), “Raindrop size distributions and radar reflectivity-rain rate relationships for radar hydrology”, *Hydrology and Earth System Sciences*, 5(4) (2001), pp. 615-627

Attached DVD

DVD is attached in wallet.

Contents:

- Code for analysis
- Videos and Images used for correlation
- Spreadsheets for rain, light and wind analysis
- Copy of Dissertation