

An On-Line Machine Vision Flotation Froth Analysis Platform

Charl P. Botha

October 1999

Thesis presented in partial fulfilment
of the requirements for the degree of
Master of Science
at the
University of Stellenbosch

Supervisor: **Prof. D.M. Weber**

Declaration

I, the under-signed, hereby declare that the work contained in this thesis is my own original work and has not previously in its entirety or in part been submitted at any university for a degree.

Signature :

Date : April 2, 2004

Abstract

Froth flotation is a popular process used for the extraction of certain minerals from ore. The effective automatic control of this process has great financial importance. Utilising machine vision as the measurement component of an automatic flotation control system has several advantages: it is a cost effective solution due to its non-contacting sensors and consequent low maintenance, as well as its low initial setup costs; it yields consistent and accurate parameters; and the measurements it makes are available immediately, thus enabling real-time control and minimising the effects of process disturbances rapidly and efficiently. In addition, machine vision can be used as a research tool for investigating froth flotation, a process which is still not fully understood.

This work documents the development of an on-line machine vision flotation froth analysis platform that is in use in the industry as the measurement component of an advanced flotation control system and in the academic world as a powerful research tool. Two additional froth analysis applications, namely an embedded froth mobility sensor for use in the flotation industry and a continuous froth video sampling and analysis research platform, are developed in the course of this work.

An investigation into continuous video sampling and analysis of flotation froth is performed for the first time and indicates that this mode of froth image acquisition yields more information than traditional froth imaging methods. Recommendations are made with respect to the practical application of continuous froth video sampling and analysis.

Opsomming

Skuim-flotasie is 'n gewilde proses vir die ontginning van sekere minerale uit erts. Die effektiewe outomatiese beheer van flotasie dra groot finansiële impak. Daar is verskeie voordele verbonde aan die gebruik van rekenaar-visie as die sensor-komponent van so 'n outomatiese flotasie beheerstelsel. Dit is 'n koste-effektiewe oplossing vanweë die feit dat die sensors nie in aanraking is met die flotasie nie en dus min onderhoud vereis. Boonop is die aanvanklike opstellingskoste laag. 'n Masjien-visie stelsel lewer akkurate en betroubare metings en dit kan hierdie metings baie vinnig beskikbaar maak wat intydse beheer van 'n flotasie-aanleg vergemaklik. Daar kan dan baie vinnig op proses-versteurings gereageer word. Masjien-visie is ook waardevol vir navorsing oor flotasie, 'n proses wat nog nie heeltemal verstaan word nie.

'n Praktiese masjien-visie gebaseerde flotasie-skuim analise platform word ontwikkel. Hierdie sagteware word gebruik in die flotasie-industrie as die sensor-komponent van 'n gevorderde flotasie beheerstelsel asook as 'n kragtige navorsings-hulpmiddel in die akademiese wêreld. Twee bykomende toepassings word ontwikkel, naamlik 'n intydse flotasie-skuim mobiliteitsensor vir gebruik in die flotasie-industrie en 'n navorsingsplatform vir die insameling en analise van kontinue flotasie-skuim video-data.

Ondersoek word ingestel na die analise van kontinue flotasie-skuim video-data vir die eerste keer in flotasie-skuim geskiedenis, en die resultate dui daarop dat hierdie metode van data-insameling meer informasie oor die skuim inhoud as tradisionele insamelings- en analise metodes. Daar word aanbevelings gemaak omtrent die praktiese toepassing van kontinue flotasie-skuim video monsterring.

Acknowledgements

Dave Weber, my advisor for the past three years, has been a good friend and mentor and has taught me everything that I know about research and coffee appreciation.

The DSP-Lab has been my research home for the past three years; the knowledge and experience I have gained from interacting with its inhabitants are considerable. It has also been a brilliant supplier of good strong coffee.

The creators of the free software that I depend on, such as GNU/Linux, CVS, LaTeX and my all-time favourite editor, JED, deserve a lot of gratitude.

Marc van Olst and Francois du Plessis have been good friends and colleagues in my interaction with Crusader Systems.

Gerhard “Goof” Esterhuizen and Mark Stoneman have taken the DSP-Group’s network to new heights of efficiency and functionality and have made sure that I’ve always had machines and good operating systems to work on.

All my friends make life a scream. Thanks guys.

My mother, Barry, Hermann and Etienne are very dear to me, and they’ve played a great part in my education as human being.

Stella has redefined the term “supportive” during times when normal humans would have given up. She is an amazing person, and my eternal love and gratitude belong to her.

Contents

1	Introduction	1
1.1	Background and definitions	1
1.1.1	The Flotation Process	1
1.1.2	Effective Control of the Flotation process	2
1.1.3	Computer Vision as Sensor Component	3
1.1.4	The Froth Surface Appearance and Process Conditions	3
1.1.4.1	Characteristics of surface appearance	4
1.1.4.2	Froth types	5
1.2	Previous Work	5
1.3	Overview of this work	7
1.4	Contribution	8
1.5	Thesis Organisation	9
2	Image Processing	10
2.1	Static Analysis	10
2.1.1	Bubble segmentation	10
2.1.1.1	Marker extraction	11
2.1.1.2	Homotopic modification	14

2.1.1.3	Watershed Segmentation	16
2.1.1.4	Post-segmentation Over-segmentation Correction	18
2.1.2	Other textural properties	21
2.1.3	Fine-froth Segmentation	22
2.1.3.1	Marker Bubble Area Ratio Threshold	22
2.1.3.2	Moving window SNE	23
2.2	Dynamic Analysis	24
2.2.1	Per-bubble motion analysis	25
2.2.2	Region-based motion analysis	25
2.3	Conclusions	27
3	Parameter Extraction	29
3.1	Pre-processing of image-processing information	29
3.2	Extractable parameters	30
3.2.1	Area/Size	30
3.2.2	Bubble-Count	31
3.2.3	Bubble-based Motion	31
3.2.4	Colour/Intensity Information	31
3.2.5	Ellipticity	32
3.2.6	Fine froth/Valid Froth Ratio	33
3.2.7	Orientation	33
3.2.8	Perimeter	33
3.2.9	Region-based Motion	34
3.2.10	SNE	34

3.3	Statistics	34
3.3.1	Linear Statistics	37
3.3.2	Circular Statistics	37
3.4	Histograms	38
3.4.1	Histogram Accumulation	39
4	Practical Implementation	40
4.1	Froth Flotation Plant Basics	40
4.2	Fluxar Flotation Control System	41
4.3	Implementations of Froth Surface Analysis	43
4.3.1	Analysis kernel	44
4.3.2	Fluxar-SP	44
4.3.3	Fluxar-Speed	48
4.3.4	Continuous sampling platform	48
5	Performance and Impact	51
5.1	On Algorithm Validation	51
5.1.1	Experiments with synthesised motion	51
5.1.2	Qualitative judgement	53
5.2	Continuous Sampling Experiment	53
5.2.1	Continuous Video Static Analysis	54
5.2.2	A Useful Filter for Froth Parameters	57
5.2.3	Continuous Video Dynamic Analysis	59
5.2.4	Conclusions	62
5.3	Performance and Impact on Flotation Plants	65

6	Conclusions	67
6.1	Summary of Work	67
6.2	Contribution	68
6.3	Future Work	69
A	Circular Statistics	74
A.1	Selected statistic measures	74
A.1.1	Mean Direction	74
A.1.2	Median Direction	75
A.1.3	Circular Variance and Standard Deviation	75
A.2	Angular Distributions on the range $(0, 2\pi/l)$	77

List of Figures

2.1	Image illustrating the over-segmentation that takes place if the marker extraction and homotopic transform is not performed before the watershed segmentation. The top image is the source image, the second is an un-preprocessed image that has been run through the watershed, and the third image has been correctly pre-processed and segmented.	12
2.2	Illustration of binary reconstruction. In the left image, the dark grey regions represent the markers and the lighter grey regions the mask. Only connected components of the mask which contain at least a pixel of a marker are kept and the rest are eliminated. The result is shown in the right-hand image.	13
2.3	Greyscale reconstruction. Analogously to binary reconstruction, the marker image marks the parts of the mask that should be kept.	14
2.4	H-dome extraction. A constant h is subtracted from an image I to yield image $I - h$. The image I is then reconstructed from $I - h$, and the result of this reconstruction is subtracted from the original image thus yielding the h-dome image.	15
2.5	Marker extraction. Note indicated position has two markers instead of one, and this causes incorrect segmentation as can be seen in the last image in figure 2.1.	15
2.6	Illustration of the homotopic modification. The marked minimum is enforced by extending it downwards to the minimum possible value and non-marked minima are filled.	17
2.7	Watershed terminology	18
2.8	Sample watershed segmentation of flotation froth	19
2.9	Oversegmentation due to extra markers	20

2.10	Results after oversegmentation correction to be compared with figure 2.1 (uncorrected segmented image at bottom) and figure 2.1.1.1 (marker image pointing out anomalous marker).	21
2.11	Images illustrating working of marker bubble area ratio threshold fine froth detection with MBRT parameter set to 0.004.	23
2.12	Results of moving window SNE fine froth segmentation	24
2.13	Sample results of per-bubble motion analysis	26
2.14	Output of the region-based motion estimation	27
3.1	Froth image, image of top 2.5% of intensity values and histogram of froth image.	36
3.2	Example of area histogram	38
4.1	Diagram of simple froth flotation circuit	41
4.2	Camera setup	42
4.3	Diagram of Fluxar Flotation Control System	43
4.4	Screen-shot of Fluxar-SP	45
4.5	Screen-shot of sample Fluxar-SP image display window. From left to right: original, bubble segmentation, region-based motion, bubble-based motion.	45
4.6	Screen-shot of Fluxar-SP database viewer with sample parameters	46
4.7	Screen-shot of Fluxar-SP Global Analysis Options dialogue	47
4.8	Screen-shot of Fluxar-SP Algorithm Options dialogue	47
4.9	Photo showing FluxarSpeed prototype internal layout	49
5.1	Small section of number of bubbles plot for 25 image per second sample set of laboratory flotation cell froth illustrating number of bubbles sensitivity.	54
5.2	Segmented froth images corresponding to graph in figure 5.1. Images 144–149 left to right, top to bottom.	55

5.3	Section of average vertical pixel motion as calculated by region-based motion calculation for 25 image per second sample set of laboratory froth. Compare with the images in figure 5.2.	56
5.4	Section of ellipticity figure plot for 25 image per second sample set of laboratory froth. See figure 5.2.	57
5.5	Plot of average bubble area/size over 400 images representing 16 seconds of on-plant tank cell froth digitised at 25 images per second.	58
5.6	Average bubble size signal filtered with α -trimmed mean filter with $N = 10$ and $\alpha = 0.2$	59
5.7	Input and segmented froth images corresponding to graph in figure 5.6. Images 117, 132 and 152 from top to bottom.	60
5.8	Plot of average horizontal pixel motion as calculated by the region-based motion estimation over 400 images sampled at 25 images per second from on-plant scavenger cell. The darker line is a version smoothed with an α -trimmed means filter with $\alpha = 0.2$ and $N = 10$	61
5.9	Plot of average horizontal pixel motion per image standard deviation as calculated by the region-based motion estimation over 400 images sampled at 25 images per second from on-plant scavenger cell. The darker line is a version smoothed with an α -trimmed means filter with $\alpha = 0.2$ and $N = 10$	62
5.10	Section of plot of average vertical pixel motion for 400 images sampled at 25 images per second from laboratory flotation cell.	63
5.11	Section of plot of average vertical pixel motion per image standard deviation for 400 images sampled at 25 images per second from laboratory flotation cell. Corresponds with figure 5.10	63
5.12	Images with motion vectors corresponding with positions 357, 360 and 363 in graphs in figures 5.10 and 5.11.	64
A.1	Illustration of median direction	75

List of Tables

1.1	Effect of process condition changes on bubble size	4
2.1	Average execution times for various algorithm groups on tank cell and scavenger cell froth.	28
3.1	Extractable parameters and available post-processing	30
3.2	Pre-processing of parameter arrays	37
5.1	Mean square pixel errors for horizontal, vertical and Euclidean motion calculated over 100 test images.	52

Keywords

image processing
machine vision
froth flotation
process control
sensor platform
greyscale morphology
motion analysis
mineral processing

Acronyms

SE	Structuring Element
NGLDM	Neighbouring grey level dependence matrix
SGLDM	Spatial grey level dependence matrix
BMA	Block Matching Algorithm
SNE	Small Number Element
LNE	Large Number Element
NNU	Number Non-uniformity
SM	Second Moment
ENT	Entropy
MCC	Maximum Cross-Correlation
SCADA	Simultaneous Control and Data Acquisition

Chapter 1

Introduction

1.1 Background and definitions

1.1.1 The Flotation Process

Froth flotation is a very widely used process for extracting certain minerals from ore. It is based on the utilisation of differences in surface properties of particles of various minerals. Mineral-rich ore is milled to a fineness that liberates most of the valuable particles and that will result in particles that are small enough to stay attached to rising air-bubbles in the pulp [1]. This milled ore is added to water. Reagents that modify the surface hydrophobicity of the particles are added to the mixture and air is bubbled through the pulp. Hydrophobic particles attach to bubbles and float to the surface where they form the froth layer and can be easily skimmed off, whereas the other particles remain in the body of the pulp [2].

In direct froth flotation the valuable mineral particles are made to be hydrophobic so that they attach to air bubbles when they collide with them in an agitated mixture of milled ore and water. Other particles should ideally remain (or be made) hydrophilic so that they remain in the pulp.

Some of the reagents which are utilised in order to realise the flotation process are [3]:

frother Enables water to form a stable froth.

collector Adsorbed on mineral surfaces rendering them hydrophobic. Facilitates bubble attachment.

activator Induces flotation in the presence of a collector which otherwise has no effect on the target mineral.

depressant Prevents collector from functioning on a mineral.

This process of extracting minerals, known as froth flotation, is used to extract various minerals including coal, copper, tin and platinum.

1.1.2 Effective Control of the Flotation process

Flotation is a very effective process for extracting minerals from very large amounts of ore and is relatively easy to implement in the extraction industry. It therefore has great financial importance [4]. Consequently, effective control of this process, resulting in better recovery and grade of mineral, is of great significance.

In most flotation plants, it is customary for plant operators to monitor the froth surface visually and make adjustments to process parameters based on their interpretation of the surface appearance and their idea of the appearance of an ideal froth. Controllable parameters include frother, depressant and collector addition rate, aeration rate, pulp level and froth height [1].

However, understanding the flotation process is extremely difficult [5] and achieving optimal control is often not possible for the human operators. There is inconsistency among operators in the interpretation of froth appearance, partly due to the fact that similar froth appearances can be caused by different sets of metallurgical parameters and the fact that there is no formal source of operator knowledge. Any combination of human error, lack of training and inadequate experience can cause suboptimal control of the plant [6].

From the above, it can be seen that there is a need for an accurate and consistent automatic control system that can assist the froth operators in their work. This system would perform monitoring utilising some kind of automatic measurement device and effect control autonomously. Such a system has other advantages as well:

1. Around the clock plant-monitoring and -control.
2. Remote monitoring of a plant where the control room is far from the plant. Some modern plants consist of several large cells ($> 100m^3$) and these are placed far away from the control room for safety-reasons, amongst others. Monitoring these plants with human operators is difficult.
3. Analysis and control actions are always consistent.
4. Concurrent monitoring of several locations in the process can be performed.
5. Sensors can be installed in ideal positions (i.e. where operators can't always reach).
6. Response on measurement is immediate.

7. All acquired measurement data are accurately and meticulously recorded.
8. Automatic measurement is potentially more accurate than human observation.

1.1.3 Computer Vision as Sensor Component

In section 1.1.4 the close relationship between the froth surface appearance and the flotation process conditions will be discussed. In addition to this relationship, there are several good arguments for having a computer vision system as the sensor component of an automatic flotation process control system, as opposed to alternative measurement methods such as x-ray fluorescence and even titration.

1. Traditional control methods employ froth operators that visually observe the froth and perform control based on their observations. A vast amount of accumulated expert knowledge is available and can be combined in such a vision system. This is one of the most compelling reasons.
2. The system will always yield consistent parameters, i.e. if it is shown exactly the same froth at two different times, it will return the exact same parameters
3. Cameras are non-contacting sensors. This makes for a low-maintenance sensing platform in a potentially corrosive environment (the froth).
4. Measurements can be made much more quickly than human operators are able to, and consequently necessary corrective measures can be quickly performed.
5. X-ray fluorescence measurements, another method of obtaining information about the froth-process, can be inaccurate [7]. These inaccuracies commonly occur at low platinum particle concentrations in extraction plants.

1.1.4 The Froth Surface Appearance and Process Conditions

It is a well-established fact in literature [7,4,8,9] and in practice that the visual appearance of the flotation froth surface is closely related to the state of the process, so that by monitoring and interpreting the surface appearance, much information about the internal process conditions can be gathered. In fact, it is common practice in flotation plants that plant operators visually monitor the froth surface and make process control decisions based on these visual observations.

Process condition/event	Bubble size	Note
too high mineral content	large	no windows; froth stable, viscous
very low mineral content	very small	watery froth
strong frother	small	
weak frother	large	
high froth viscosity	large	
increase in frother concentration	decrease	
increase in pH	increase	
presence of ions in solution	decrease	better collection
temperature increase	decrease	

Table 1.1: Effect of process condition changes on bubble size

1.1.4.1 Characteristics of surface appearance

Bubble size

The bubble size distribution in the pulp has great effect on the probability of collision between mineral-rich particles and bubbles, as well as the adhesion of these particles to the bubbles. Bubble size on the surface and bubble size in the pulp are assumed to be related. Bubble size is also related to process variables such as the aeration rate, fluctuations in pH, the addition of reagents and the nature and amount of the minerals in the surface froth. Bubble size is one of the more important froth parameters.

Table 1.1 shows the relationship between bubble size and selected process conditions. In addition to the factors mentioned in the table, it has been established that bubble size should be adjusted to compensate for different particle sizes [7]. Specifically, it is suggested that decreasing bubble size can improve flotation performance with small particles drastically. Air flow rate has a marked effect on the grade and recovery of mineral and as there is a relation between air flow rate and bubble size, bubble size is significant in this respect as well.

Bubble shape

Polyhedral bubble shapes occur as a result of bubbles being packed closely together in a dry froth. This phenomenon indicates good drainage and a high grade. Spherical bubbles on the other hand imply that there is more water between the bubbles, and therefore the froth is runny and has low mineral content. Elliptic shapes are caused by high froth viscosity and stability.

Froth viscosity and mobility

Viscosity is influenced by the amount of hydrophobic particles present in the froth. The viscosity and mobility are inter-dependent and play an important part in flotation performance [8].

Froth stability

When a froth is too stable, it could mean that the depressant flow is too low and that the bubbles contain too much gangue (waste minerals). This obviously results in a low grade concentrate. If a froth is unstable, its mineral-laden bubbles burst when they go over the cell-weir.

1.1.4.2 Froth types

Three primary froth types are identified by Moolman et al. in [7, 4]. Obviously the specific appearances of the froth types differ for different flotation setups.

Ideal froth The bubbles are closely packed with bubble windows, and the froth is not too viscous nor too runny. The bubbles are larger than those of the runny froth and smaller than those of the sticky froth.

Runny froth The mineral load is low and the froth is watery, unstable and very mobile. The bubbles are small.

Sticky froth The bubbles are large and can be elliptical. The froth is very stable and viscous.

1.2 Previous Work

In this section, we look primarily at work pertinent to the automatic control of the froth flotation process utilising computational vision as the sensor component as well as just computational vision froth surface analysis.

In 1992, Symonds and De Jager [10] first proposed the segmentation of images of the froth structures into bubbles using morphological methods. These bubbles could then be counted and analysed to characterise the segmented froth images. The morphological segmentation algorithm that they proposed was a residual edge detector utilising a spherical structuring element (SE). By opening and closing the image with a spherical SE (“rolling ball algorithm”) and then subtracting the opened and closed images from the original image, a peak image and a valley image could be

obtained. These attempts, although promising, weren't entirely successful due to the complexity involved in trying to identify individual elements (bubbles) in the froth structure [8].

Moolman et al. took a different approach in [11] and used second order statistics to characterise the textural properties of the surface froth. These textural properties were used as input features to neural network classifiers. Using these techniques, they could accurately classify froth type and detect froth mobility. Their method was refined in [8] to increase its classification accuracy. The textural features that they extracted were numerical features on the neighbouring grey level dependence matrix (NGLDM) and spatial grey level dependence matrix (SGLDM) based, both second-order grey level statistics [12].

Eksteen's work [1] utilised the image processing software and algorithms as developed by Moolman et. al, but focused more on providing a knowledge base for the research and development of vision-based automatic froth control systems from a chemical engineering viewpoint. This work concluded that a vision-based system for automatic flotation control was a commercially viable venture and put forth the parameters for this viability. As a general overview of the factors pertinent to the development of a complete vision-based automatic flotation control system, this is a valuable work. From an image-processing point of view, it contains a more detailed discussion of the methods employed by Moolman et al.

Note that by this time, there were two distinct approaches in the froth surface visual analysis business. One approach attempted to analyse the froth structure in a hierarchical and direct fashion by attempting an accurate segmentation of the froth surface into bubbles, and then calculating higher level parameters using this segmentation [10]. The other approach related distinct froth types to textural properties and could indirectly derive other froth surface visual parameters from texture.

The first successful attempt at full bubble segmentation was made in 1997 by N. Sadr-Kazemi and J.J. Cilliers of UMIST [13]. In this research, the morphological watershed algorithm was used to segment froth images into bubbles. It is also stated that "there have been few attempts at estimating the complete bubble size distribution and relating that to the flotation performance".

The dynamic (motion) analysis of flotation froth is also significant as it is interdependent with the froth viscosity. Motion blur caused by slightly longer than normal exposures on video capturing equipment was used by Moolman et al. [9] and Eksteen [1] to calculate froth mobility. They found that the NGLDM second moment as calculated on one of these blurred images gave an accurate measure of froth speed. This conforms with their paradigm of classifying froth types according to surface texture. Other work on calculating froth mobility [14] used standard block matching algorithm (BMA) based methods and compared these with bubble-based motion. BMA based methods calculate motion by attempting to match small regions of a source image with corresponding regions on a target image, where the source and target image are images of the same scene taken temporally apart. The bubble-based motion estimation was found to be more suitable, contrary to our results.

A complete vision-based expert system for flotation froth control is documented in [15]. Bubble-segmentation is performed, but the description of the method used is quite vague, although it does seem related to the watershed. This system, although commercially marketed, does not seem very successful.

JKFrothCam is a commercially available system for the on-line monitoring of flotation froth that was developed at the Julius Kruttschnitt Mineral Research Centre (JKMRC) in Australia [16]. It captures, digitises and analyses video from cameras monitoring the flotation froth surface. It is able to calculate three surface characteristics, namely average bubble size, froth texture and froth velocity.

For the froth texture, a texture unit number is generated at each point by performing a certain calculation (measuring the difference between neighbouring pixels) within a 3×3 pixel matrix at that pixel. A histogram, or texture spectrum, of number of occurrences vs. texture unit number is then generated. The height of the middle peak in this spectrum is used to classify bubble size, whilst the whole spectrum is compared to pre-calibrated representative froth spectra to determine the froth type.

The velocity is calculated using what the JKMRC calls “pixel tracing”. Judging by available documentation this is a form of the BMA for motion estimation.

JKFrothCam is able to estimate certain metallurgical properties of the froth, such as ash concentration, using a simple linear model with the three extracted parameters (type, bubble size, velocity) as variables. The model parameters are regularly calibrated using a self-tuning Kalman filter.

1.3 Overview of this work

This work documents the development, implementation and investigation of a complete set of flotation froth surface computational vision analysis algorithms. These algorithms are able to perform a complete static and dynamic analysis on digital images of the froth surface.

The static analysis includes a full bubble segmentation where the perimeter of each bubble is determined using watershed morphology. In addition to the number of bubbles, parameters like per-bubble area, perimeter length, orientation angle and ellipticity are calculated from the segmentation results. As part of further post-processing, histograms and statistic values are calculated, taking special care with angular data.

Textural parameters are also calculated as part of the static analysis and can be used as an auxiliary indication of parameters like bubble-size. In addition, areas that contain fine froth, i.e. froth that is too fine to segment into distinct bubbles, are tagged as such. This is a very important step,

as big areas of fine froth are meaningful characteristics. A full colour and intensity analysis is also performed.

The dynamic analysis yields parameters that are an indication of froth mobility. It is performed on two consecutive images. A region of interest-based motion analysis is performed, and this yields motion vectors at regular intervals within the user-definable region of interest. This is usually done over the froth weir. A bubble-based motion analysis yields a motion vector for each bubble. These are not at regular intervals. Statistic values and histogramming can be calculated on these sets of parameters as well.

As part of this work, these algorithms have been implemented in a practical software package for integration in on-plant automatic control systems. This sensor platform, consisting of the algorithm software, a graphical user interface and other functionality, is able to acquire live video from multiple input channels and perform complete froth analysis in pseudo-realtime. The acquisition and analysis is fully configurable by the software operator utilising the user-friendly interface.

The sensor platform has been practically integrated with an automatic closed-loop control system in an on-plant environment where it performs well. The whole system, including the sensor platform, is being actively marketed. It is also being used in research environments locally and abroad.

Another implementation of the platform is an embedded solution that calculates a subset of the complete analysis. This product is in the process of going to international market.

This work also investigates the analysis of flotation froth on continuous full-rate video. In practical setups, complete analyses are performed at 30 to 120 second intervals as the time constants for most control actions on a flotation plant are quite long. There is merit in performing analysis at full video frame-rate however, i.e. 25 complete analyses per second (European standard). Experiments on using such full-rate analysis were performed and are documented in this work.

1.4 Contribution

This work has resulted in a robust sensor platform that performs both full bubble segmentation and extracts textural parameters, thus combining the two froth analysis approaches discussed in section 1.2.

Due to this combination of paradigms, the fact that it performs two distinct kinds of mobility analysis and the particular fine froth detection functionality (which has not been documented elsewhere), it is probably the most complete froth surface visual analysis package available to date. It is applicable in industrial and research environments.

In order to facilitate the development of this package, a practically re-usable froth analysis code kernel was created that can be re-deployed in any number of practical applications. One such application is the embedded froth surface mobility analyser that was developed in the course of this work.

In contrast to the long analysis times required by the bubble segmentation implementations mentioned in section 1.2, this work is able to perform full analyses in near-realtime, enabling it to be deployed on flotation cells in a closed-loop control system.

The sensor platform has been successfully integrated with a practical flotation plant control system, in which capacity it is performing very satisfactorily. It is one of only two such systems known in the world that are currently being marketed, and it far outstrips the competing sensor platform with respects to functionality.

No known prior work has investigated froth surface analysis on continuous full-rate video data. We have performed experiments in this regard and our results are documented in this work. A re-usable application that is able to gather large amounts of froth images at continuous rates and analyse these off-line was developed using the froth analysis kernel. This is a valuable research tool.

This work adds to the small formalised knowledge-base available to practitioners and researchers in the machine vision analysis of flotation froth.

1.5 Thesis Organisation

Chapter 2 documents the specific image processing algorithms and principles that were used to analyse the froth structure images. At top level, the chapter can be divided into static and dynamic analyses.

In chapter 3 we detail how the information yielded by the image processing algorithms and techniques is processed and converted to usable parameters.

A description of the implementation of the whole froth surface visual parameter extraction system is given in chapter 4. This also describes the complete control platform of which the sensor platform is an important part.

Chapter 5 presents the results obtained from both laboratory and plant operation.

Chapter 6 contains a summary of this work, the contribution therein and re-iterates the important knowledge that was gained. Possibilities for future work on this subject are also discussed.

Chapter 2

Image Processing

This chapter documents two classes of low-level image algorithms that were investigated. The first, “static analysis”, refers to algorithms that are usually performed on a single static image and render information which is pertinent to froth structure and form. The second class, “dynamic analysis”, refers to algorithms that are performed on two or more images and utilise these images to determine froth motion information.

Chapter 3 details the processing that converts the results of this low-level analysis into higher level characteristics that describe the nature of the flotation froth surface.

2.1 Static Analysis

Static analysis includes all parameters that can be determined from a single image and thus excludes any motion and other dynamic analysis. This stage implements a bubble segmentation algorithm, fine froth detection and the calculation of global second order statistical textural properties.

2.1.1 Bubble segmentation

For the purposes of the static analysis, it was decided that a full bubble segmentation would be required in addition to textural analysis, thus combining the two froth analysis approaches discussed in section 1.2. This decision was significant in the light of the fact that such a full segmentation is very expensive in terms of processing time and that one of the requirements of this work was real-time operation. None of the previous full-segmentation implementations document real-time operation.

From an accuracy perspective, extracting structural bubble characteristics such as size, shape and perimeter is better done by physically segmenting bubbles and calculating the true quantities rather than using a measure which is *indirectly* related to these quantities, such as the approach followed the JKMRC in [16] and Moolman et al. in [9] where certain textural properties were used to estimate e.g. the average bubble size. In addition, having structural information available for each and every bubble makes it possible to calculate complete parameter histograms. Whether the segmentation approach is more or less applicable to the measurement of froth for control purposes will be discussed in chapter 5.

Because of the way that watershed morphology intuitively applies to froth images as will be seen in section 2.1.1.3, it was decided to perform element marker-extraction followed by watershed segmentation as documented by Vincent [17, 18] for the general case, and as applied to flotation froth images in [13] by Sadr-Kazemi.

Vincent's implementation was found to be extremely fast and accurate, and evolved, with a few changes, into the segmentation functionality of the final product. In [13], off-the-shelf implementations of the algorithms are used. In this work, highly optimised versions from [17] and [18] were implemented and afforded on-line execution speed of the segmentations. An indication of the speed of execution is given in section 5.2.

The following sections describe the different stages of the complete bubble segmentation of a flotation froth image.

2.1.1.1 Marker extraction

In order to prevent over-segmentation by the watershed algorithm, the image that is to be segmented has to be pre-processed in a certain way. Figure 2.1 illustrates the effect of this pre-processing. Note that the second image has been severely over-segmented, whereas the third image, which has been pre-processed, can be considered to be a correct segmentation.

This pre-processing step requires a binary marker of arbitrary size for each expected element. That is, for each bubble that we expect to segment out, precisely one contiguous collection of coordinates within the area of that bubble is required. The pre-processing step, called the homotopic modification, is explained in section 2.1.1.2 and the requirement thereof by the watershed is discussed in section 2.1.1.3.

In practical setups it is found that each bubble displays a localised reflection, either due to the sun or to the spotlights that are used on industrial camera setups. These localised reflections are always available where there is one clearly dominant source of lighting approximating a point source. Symonds and De Jager used them as an indirect and proportional measure of bubble size and shape [10] whilst other work used them to extract bubble-markers [13, 15].

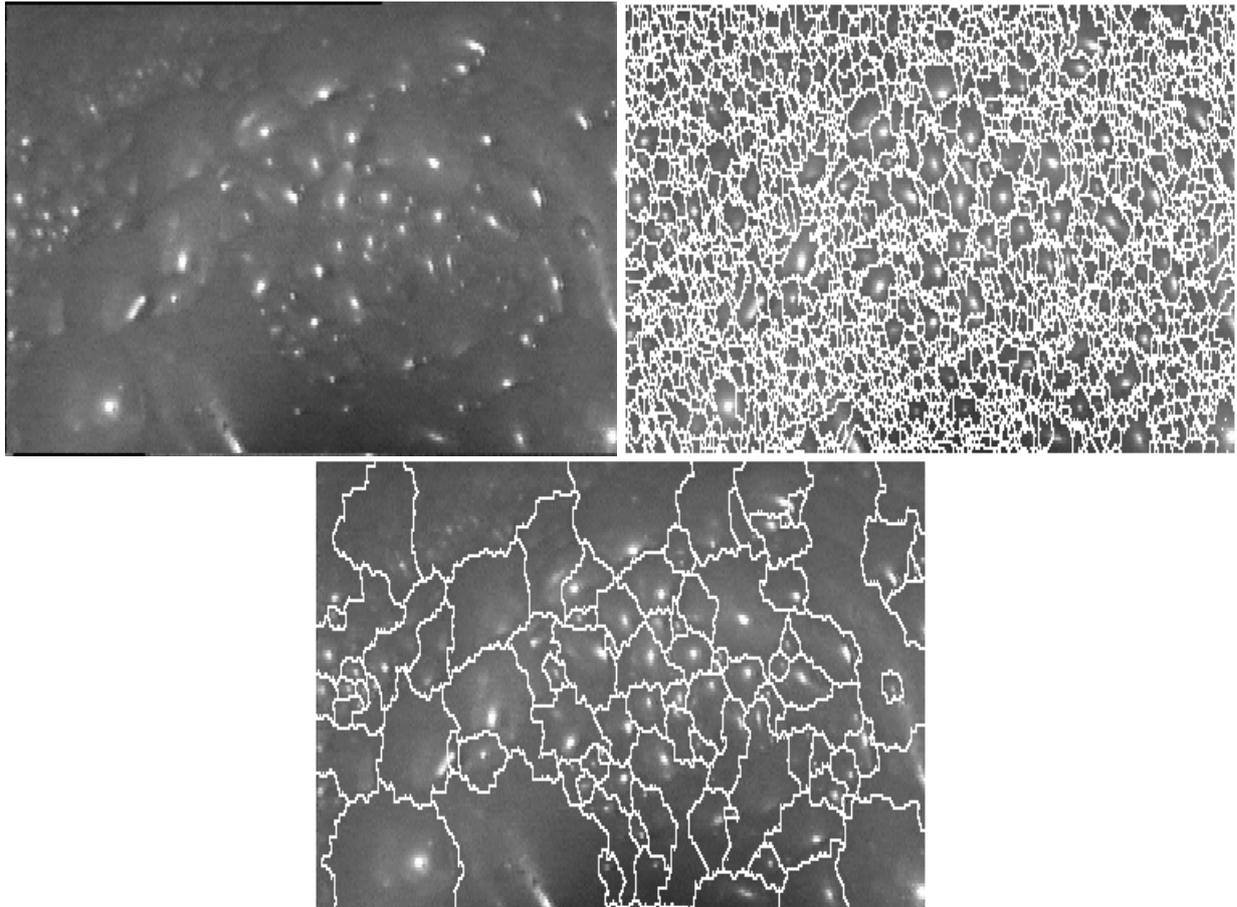


Figure 2.1: Image illustrating the over-segmentation that takes place if the marker extraction and homotopic transform is not performed before the watershed segmentation. The top image is the source image, the second is an un-preprocessed image that has been run through the watershed, and the third image has been correctly pre-processed and segmented.

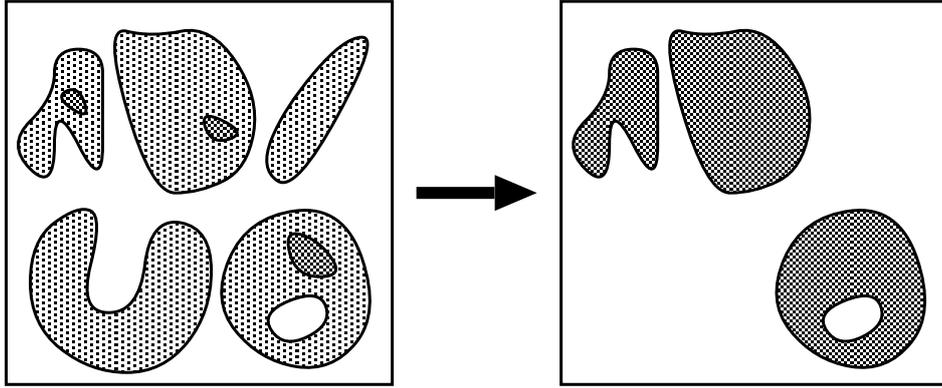


Figure 2.2: Illustration of binary reconstruction. In the left image, the dark grey regions represent the markers and the lighter grey regions the mask. Only connected components of the mask which contain at least a pixel of a marker are kept and the rest are eliminated. The result is shown in the right-hand image.

The h-dome extraction [18] was used to extract these localised reflections. This transform is able to extract light structures, i.e. regions of pixels which are brighter than their immediate surroundings, without any size or shape criterion. An h-dome D of image I is defined as a connected component of pixels such that every pixel p that is a neighbour of D satisfies $I(p) < \min\{I(q)|q \in D\}$ and $\max\{I(q)|q \in D\} - \min\{I(q)|q \in D\} < h$, i.e. all the neighbouring pixels of the h-dome, or light structure, are less bright than the light structure itself, and the maximum intensity range of the pixels in the light structure is less than h . The value of a pixel p in the h-dome D is $I(p) - \min\{I(q)|q \in D\}$, i.e. the pixel values are offset at the minimum intensity value of the h-dome.

The h-dome extraction relies on a process called grey-scale reconstruction. Before grey-scale reconstruction can be discussed, binary reconstruction has to be defined. The binary reconstruction $\rho_I(J)$ of mask image I from marker image J is the union of all the connected components of I which contain at least a pixel of J . I and J in this case are binary images that satisfy the condition $J \subseteq I$. Figure 2.2 illustrates binary reconstruction.

We now define grey-scale reconstruction. Let I and J be grey-scale images both defined on the domain D_I , taking values in the discrete set $[0, N - 1]$ and satisfying $J(p) < I(p)$ with $p \in D_I$. Pixel p is an element of D_I . The grey-scale reconstruction $\rho_I(J)$ of mask image I from marker image J is defined as:

$$\rho_I(J)(p) = \max\{k \in [0, N - 1] | p \in \rho_{T_k(I)}(T_k(J))\}$$

where $T_k(I) = \{p \in D_I | I(p) \geq k\}$ is the binary thresholded image of I at value k and so $\rho_{T_k(I)}(T_k(J))$ is the binary reconstruction of mask binary image $T_k(I)$ from binary marker image $T_k(J)$. In other words, for each pixel p , the grey-scale reconstruction is the maximum intensity value k at which a contiguous region of the binary thresholded image of I is “on” over that pixel p and is marked somewhere by the binary thresholded image of J , where both binary images have been thresh-

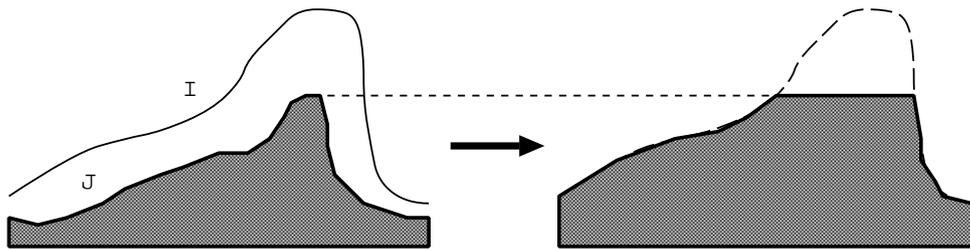


Figure 2.3: Greyscale reconstruction. Analogously to binary reconstruction, the marker image marks the parts of the mask that should be kept.

olded at intensity value k . Figure 2.3 illustrates the grey-scale reconstruction of mask I from marker J .

The grey-scale reconstruction is implemented in a highly optimised fashion using iterated grey-scale dilations and infimums (point-wise minima). Conceptually, the marker image J is grey-scale dilated by a flat structuring element B . J is now replaced by the infimum of mask image I and the result of the dilation. This is repeated until stability is reached, i.e. no change in J takes place. This corresponds to the definition in the previous paragraph, but lends itself to effective implementation and computational optimisation.

The h-dome extraction works by subtracting a constant, h , from all pixel intensities in an image I , thus yielding image $I - h$. I is then greyscale reconstructed from $I - h$ and the resulting reconstruction is subtracted from the original I . This yields an h-dome extraction of image I . This principle is illustrated in figure 2.4.

The extraction was found to be very effective and robust on froth surface images taken under various conditions, lighting and otherwise. The value of h chosen is very important but is effective within a relatively wide range. We found that for most froth images $h = 50$ yielded excellent results. As a binary marker image is required, the resulting h-dome extraction has to be thresholded. In an attempt to eliminate any undesired results (e.g. two h-domes for a single localised reflection, a possibility with certain bubble-structures as can be seen in figure 2.5) the threshold was chosen at $h/2$. Figure 2.5 shows an example of the h-dome extraction and thresholding as performed on an image of the froth in an industrial tank cell. Note how well the localised reflections are extracted, in spite of the fact that they differ in size, shape and intensity.

2.1.1.2 Homotopic modification

The homotopic modification [17, 18] utilises the extracted markers and modifies the image that is to be analysed in such a way that:

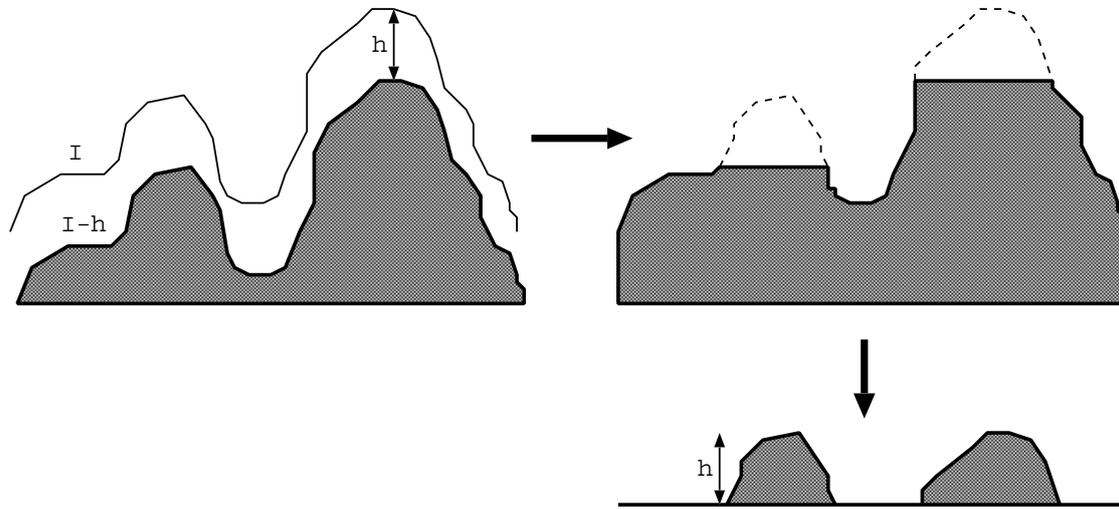


Figure 2.4: H-dome extraction. A constant h is subtracted from an image I to yield image $I-h$. The image I is then reconstructed from $I-h$, and the result of this reconstruction is subtracted from the original image thus yielding the h-dome image.

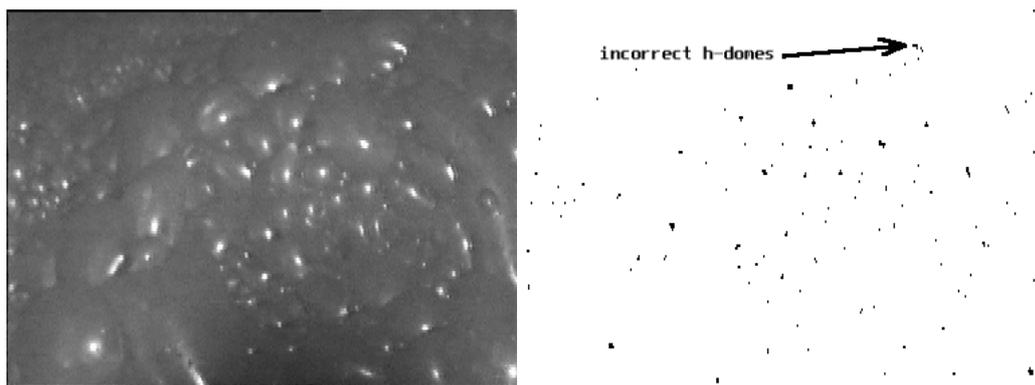


Figure 2.5: Marker extraction. Note indicated position has two markers instead of one, and this causes incorrect segmentation as can be seen in the last image in figure 2.1.

1. the image's only minima are located on the extracted markers,
2. all the other minima in the image are suppressed,
3. the most significant intensity divisions between marker positions are preserved.

These changes prevent the watershed algorithm from over-segmenting the image that is to be analysed.

In order to perform this modification the dual greyscale reconstruction $\rho_I^*(J)$ of mask I from marker J is required. The dual greyscale reconstruction is defined in terms of iterated greyscale erosions and supremums. Marker image J is eroded with flat structuring element B . J is now replaced by the supremum of the result of the erosion and mask I . This process is repeated until it stabilises, i.e. J shows no change. From this definition it is easy to see that by inverting both the marker and the mask image, then performing a standard greyscale reconstruction, and then inverting the result, one can effect dual greyscale reconstruction.

The homotopic modification $H(I)$ of image I using binary marker image M is defined as:

$$H(I) = \rho_{\min(I, \max(I) \times M)}^*(\max(I) \times M)$$

where the marker image M has zeroes where minima are to be enforced and ones otherwise.

Seeing the difference between a modified and unmodified image with the naked eye is very difficult. In general, a homotopically modified image appears “smoother” and “cleaner” (i.e. with less noise) than an unmodified one.

Using the descriptive explanation of the standard greyscale reconstruction in section 2.1.1.1 and the inversion principle discussed above, one can visualise the working of the homotopic modification and especially how it succeeds in removing non-marked local minima. Figure 2.6 illustrates the working of the homotopic transform on an inverted bubble and a marker.

2.1.1.3 Watershed Segmentation

The greyscale watershed [17, 19] is a morphological segmentation algorithm and is the cornerstone of our bubble segmentation. In order to understand the watershed, it is necessary to visualise a greyscale image as a three dimensional landscape, or a topographic relief, where a pixel intensity represents the altitude of the landscape at that Cartesian position.

This landscape consists of valleys and hills, representing the darker and brighter portions of an image respectively. If a drop of water were to fall at any position on the landscape, it would flow down and come to rest in the deepest part of the closest valley. This deepest part represents a local minimum in the image. The whole area surrounding it, from whence a drop would flow

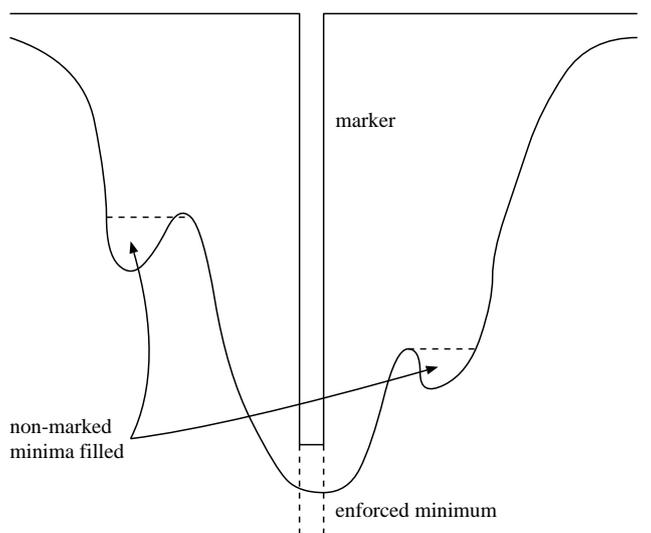


Figure 2.6: Illustration of the homotopic modification. The marked minimum is enforced by extending it downwards to the minimum possible value and non-marked minima are filled.

down into that specific local minimum, is the catchment basin of the local minimum. The borders that separate catchment basins belonging to different local minima are called watershed lines or just watersheds. Figure 2.7 illustrates the terms just defined.

The purpose of the greyscale watershed is to extract watershed lines, i.e. contiguous lines of high-intensity pixels separating local intensity minima in the image. An alternative approach to understanding the watershed that is closer to the implementation that we used is the “immersion” definition [17]. If we were to puncture the landscape at all its minima and then immerse it in a lake, the water would progressively fill all the catchment basins. At each pixel where the water from two different minima would merge, we build a dam wall. When the whole landscape has been fully immersed, the set of resultant dam walls represent the watershed lines that delineate the whole landscape into its constituent catchment basins.

The immersion definition also makes the applicability of the watershed algorithm to bubble segmentation obvious. Recall that we need to segment out each and every bubble in an image of flotation froth surface. Each bubble is structurally an irregular dome, with the lowest elevation or altitude being where bubbles meet, i.e. the domes intersect, and the highest point usually coinciding with the localised light reflection if the light source is close to the the normal axis of the froth surface. If we approximate the elevation at any point with the recorded pixel intensity at that point, which is intuitively and empirically a valid approximation, we have a digital elevation map of the froth surface. If we invert the elevation map, i.e. turn the maxima into minima and vice versa, and perform the watershed immersion on the inverted elevation map, the result is a set of watershed lines coinciding with bubble boundaries; in other words, a segmentation of the bubbles in the froth image. In practice, the result is a tessellation of the image into catchment

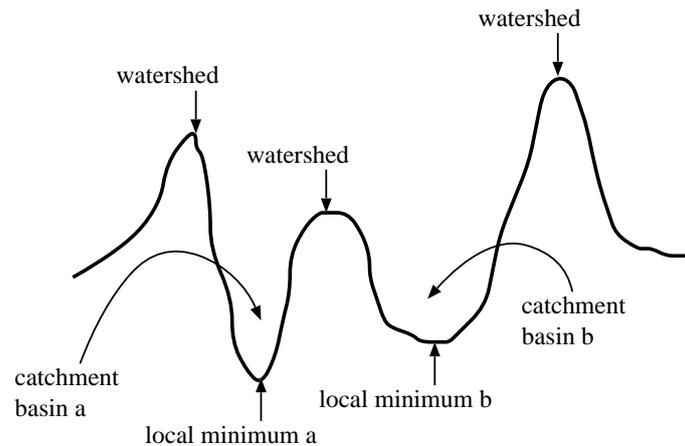


Figure 2.7: Watershed terminology

basins where each pixel is tagged depending on the catchment basin that it belongs to.

With this explanation in mind, one can also see why the homotopic modification is so important. After inversion, bubbles, being irregular domes, have numerous spurious local minima. From the immersion definition, one can see that extra watershed lines will result for all these extra local minima. The homotopic modification fills these extra local minima so that the watershed extracts only the desired bubble boundaries.

Our implementation is the highly optimised watershed algorithm from [17]. It executes the watershed by performing immersion in a very computationally efficient method. The reader is referred to the reference for details. Previous to this specific implementation of the watershed, it was not possible to perform such a segmentation in near-realtime.

Figure 2.8 shows a sample result of the watershed on a froth image that was homotopically modified beforehand. Note how well the watershed lines (indicated in black) follow the bubble boundaries.

2.1.1.4 Post-segmentation Over-segmentation Correction

As we have just seen, the homotopic modification uses a marker image to ensure that the watershed segments out only one catchment basin per bubble. This process assumes that a single marker will be extracted for each bubble. This assumption almost always holds true. There are occasions however when the irregularity of the bubble domes causes major discontinuities, e.g. darker lines, in the localised light reflections that are used as markers. See figure 2.5 for an example. These discontinuities result in more than one marker being found for that specific bubble, which in turn results in that particular bubble being over-segmented, i.e. segmented incorrectly

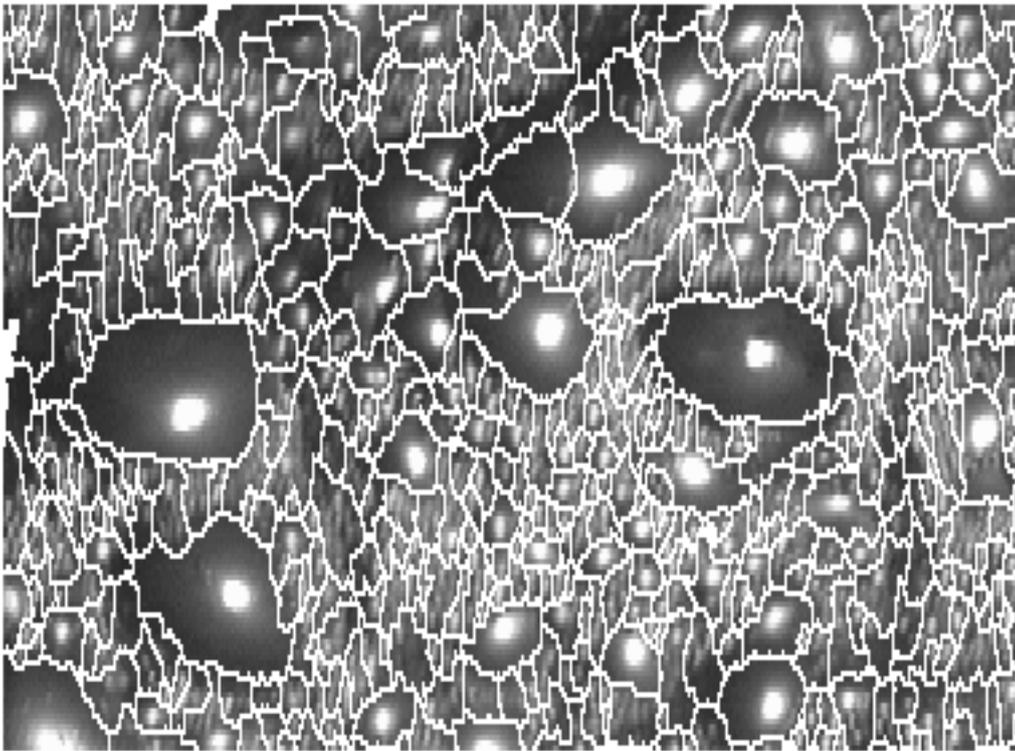


Figure 2.8: Sample watershed segmentation of flotation froth

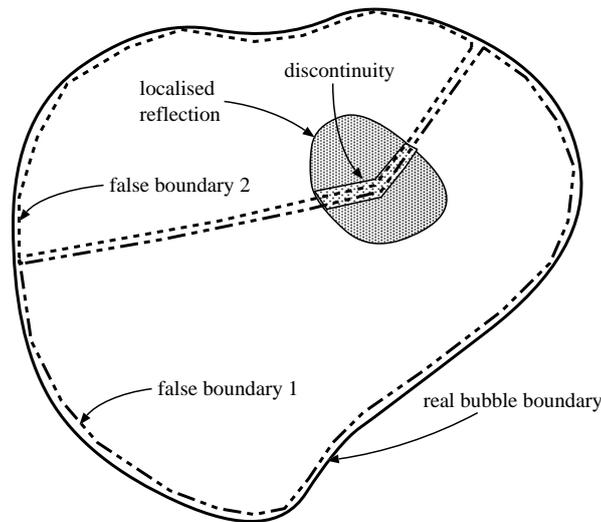


Figure 2.9: Oversegmentation due to extra markers

as two or more bubbles. Figure 2.9 illustrates this over-segmentation of a bubble (for illustrative purposes, darker regions represent higher intensity regions of the image). Note that because of the indicated darker discontinuity in the localised reflection, two markers get extracted. The marker extraction is quite robust to variations of intensity in the localised reflections but major discontinuities will cause errors.

The post-segmentation oversegmentation correction algorithm¹ is a rule-based method of remedying oversegmentation of bubbles. Initially, all boundaries, i.e. pixels on the borders of catchment basins, are calculated from the watershed tessellation. Each boundary is examined in turn for any pixels which have a grey intensity value above a certain relatively high threshold. We found that for froth images that have been normalised to the full available greyscale and for the parameters we use for the h-dome extraction, a threshold of 210 yields good results. If such a pixel is found, it is assumed that it sits on a false boundary, i.e. a boundary intersecting a localised light reflection. The first neighbouring pixel that belongs to a different catchment basin is then found. These two catchment basins are then merged and a new boundary is formed around the combination. The two old boundaries are deleted. In this way, if a bubble has been incorrectly segmented as two distinct bubbles, the two bubbles will be merged, thus rendering a correct segmentation result. Figure 2.10 an image that has been oversegmentation corrected. Compare with the uncorrected segmented image at the bottom of figure 2.1 and the marker image in figure 2.1.1.1 that points out an anomalous marker.

¹Algorithm by David M. Weber

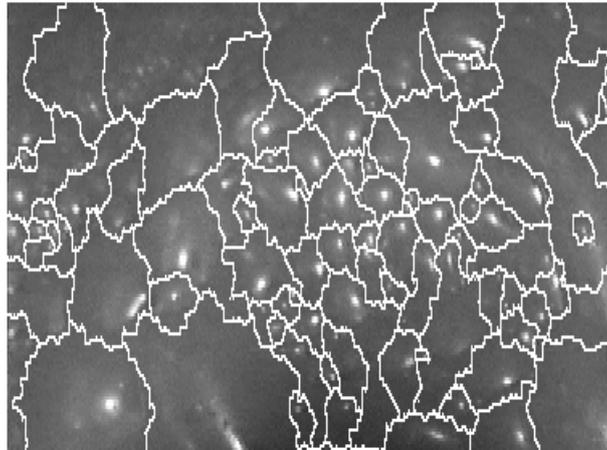


Figure 2.10: Results after oversegmentation correction to be compared with figure 2.1 (uncorrected segmented image at bottom) and figure 2.1.1.1 (marker image pointing out anomalous marker).

2.1.2 Other textural properties

As noted in section 1.2, Moolman et al. based their froth classification efforts on the extraction of SGLD and NGLD matrices and the calculation of features on these matrices. Some of this functionality has been included in this work for reasons which will be discussed shortly.

The Neighbouring Grey Level Dependence Matrix (NGLDM) [12] as calculated on an image or image region contains information about the relationship between each of the pixels and all its neighbouring pixels at a certain distance d . The matrix is represented as a two-dimensional array Q of size Ng rows by Nr columns, where Ng is the number of possible grey levels and Nr is the number of possible neighbours to a pixel in an image. The Q matrix is computed for a specific positive integer value of distance d and difference a . For every pixel, the number of times that its grey value and a neighbouring pixel's grey value differ by a or less are counted. The value at row i and column j in the Q matrix is incremented, where i is the grey value of the pixel under examination, and j is the value itself.

This matrix relates to the texture information in an image or image region. It is invariant to image rotation as well as to linear grey level transformations or shifts of the grey level spectrum. For example, if the brightness of an image were to be increased without resulting in any values being truncated, the NGLDM would look exactly the same as in the darker case.

Scalar image features can be calculated on the Q matrix, such as the Small Number Emphasis (SNE), Large Number Emphasis (LNE), Number Non-uniformity (NNU), Second Moment (SM) and Entropy (ENT). Each of these scalars indicates some qualitative measure. The SNE is a measure of the fineness of the image, the LNE is a measure of the coarseness of the image, the SM is related to the homogeneity of the image. The NNU and ENT are also related to the

coarseness of the image. In this work, we make use of only the Small Number Element.

The SNE is defined as:

$$SNE = \sum_k \sum_s [Q(k,s)/s^2]/R$$

where $R = \sum Q$ is a normalisation factor. A fine image, i.e. consisting of many grey level variations per area, will have more entries in the lower values of the Q matrix columns (representing less neighbouring pixels having similar grey values) making $Q(k,s)/s^2$ large for small s and resulting in a larger SNE. Conversely, a smooth image will have a much smaller SNE.

Initially, the SNE was used as a feature input to neural-network based froth type classifiers. In recent times however, it has become useful as a secondary indication of bubble size. In fact, in the automatic flotation control systems designed by Crusader Systems, it has become an important parameter of process stability in its own right². The SNE remains very stable for a particular froth until a process disturbance occurs, in which case it deviates noticeably.

Because the SNE is an accumulated measure of fineness for a whole image calculated from an already accumulative dependence matrix, there are billions of image configurations that will give similar SNEs. It is not a very sensitive parameter. It can afford a high degree of stability over time by trading off on accuracy. For the effective control of flotation however, it is a valuable, although coarse-grained, observation.

2.1.3 Fine-froth Segmentation

Certain areas of the flotation froth surface consist of bubbles that are too small to segment out. It is important that these areas are detected so that they can be compensated for. In addition, knowing which percentage of the froth surface is covered by fine froth is a valuable parameter, as it is an indication of the average bubble size when the per-bubble segmentation is not able to segment the whole image. This is a very difficult problem, as the froth textures are very complex.

In cases where fine froth areas are present, the watershed segmentation will segment out these fine froth areas as large bubbles. This is because the occurrence of relatively bright localised reflections in fine froth is scarce, and this results in a small number of markers being extracted for relatively large areas of fine froth.

2.1.3.1 Marker Bubble Area Ratio Threshold

This method makes use of the behaviour of the bubble-segmentation documented above. The segmentation algorithm segments out large regions of fine froth as single bubbles. This implies

²The author has investigated the use of this work on a South African platinum-extraction plant.

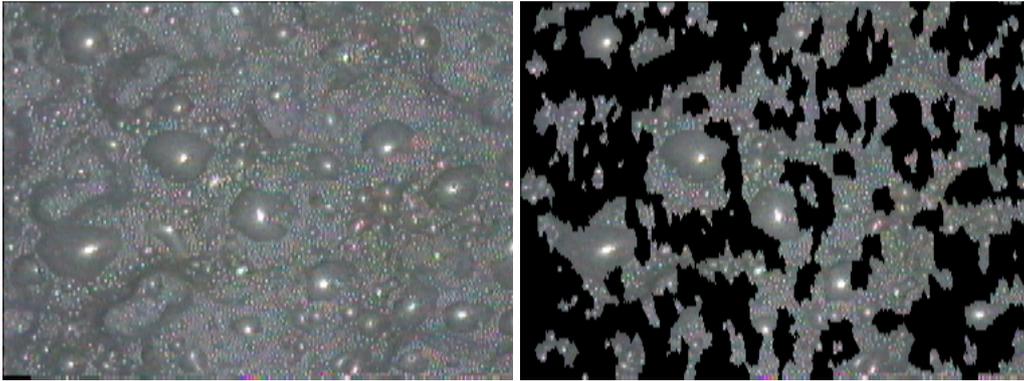


Figure 2.11: Images illustrating working of marker bubble area ratio threshold fine froth detection with MBRT parameter set to 0.004.

that the ratio of the area of the extracted marker for a specific fine froth region to the area of the fine froth region itself is considerably lower than for normal segmentable bubbles, where the localised reflection and, by implication, the extracted marker are usually relatively large.

By calculating this ratio for each and every segmented area and by comparing it to a certain threshold, areas with a lower ratio than the threshold can be tagged as fine froth. The threshold depends on lighting, froth type and froth cell type, among other factors. It is left as an operator-adjustable parameter.

This method is not absolutely accurate, but it has been empirically found that it improves the results of the static analysis. The threshold parameter determines the sensitivity of the fine froth detection, and can be optimised either for minimum false alarms or for highest detection rate. In addition, because the marker size and corresponding bubble size are already available, this segmentation is very fast. This method is used in the froth analysis product.

Figure 2.11 shows the result of a marker bubble area ratio threshold (MBRT) based fine froth detection.

2.1.3.2 Moving window SNE

Another approach that can be taken for fine froth segmentation is the more obvious one of calculating a region-based texture measure and classifying areas on the basis of this measure. For the purpose of this application, an un-supervised form of classification is required.

A moving window SNE approach was investigated and yielded promising results. In this approach a window of certain size is moved over the image in question and the SNE is calculated at every pixel position. In this way, a feature image is generated. By performing simple thresh-

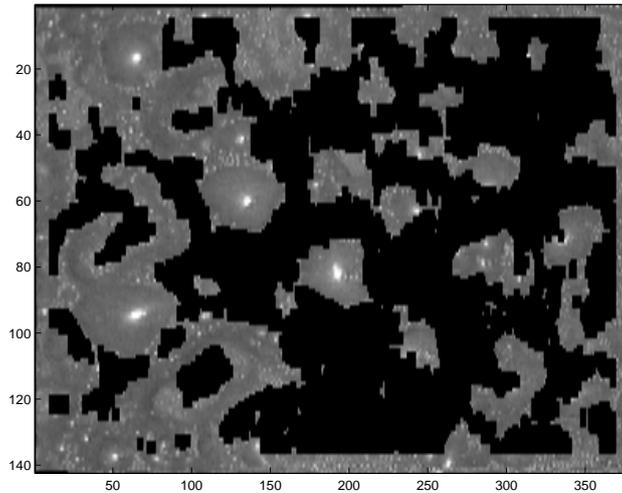


Figure 2.12: Results of moving window SNE fine froth segmentation

olding on the feature image, fine froth areas can be segmented.

Figure 2.12 shows the results for such a segmentation on the image shown in figure 2.11. This sample was generated with a window size of 20×10 , the SNE feature image was thresholded at 0.85 and the result was morphologically opened [19] in order to remove noise. Compare these results with those in figure 2.11. This method has the advantage that it is not limited by the bubble segmentation. Keep in mind that the MBRT method can only tag whole segmented regions as fine froth, whereas this method can do this on a pixel by pixel basis.

Unfortunately, this method is very computationally expensive and its parameters are much more sensitive and difficult to adjust than those of the MBRT. Because of this, it is not applicable in a practical froth analysis product.

2.2 Dynamic Analysis

The mobility of the froth is a significant parameter as it is related to the viscosity and the amount of hydrophobic particles in the froth. As such, the mobility is related to frother addition rate, depressant flow rate, activator flow rate, pulp level and density, and aeration rate.

2.2.1 Per-bubble motion analysis

This particular algorithm was first documented in [20]. It works by analysing the motion undergone by the localised reflections on the bubbles. In order to do this, the marker extraction discussed in section 2.1.1.1 is used.

As input, two images of the same froth scene taken a short time apart are used. We shall refer to these as the source image and the target image. These two images are added to form a new sum-image. If the time interval between the two images is not too long, this will result in an image with enlarged localised reflections, each of which encompasses its two constituent localised reflections in the source and target image. A slightly longer time interval will cause enlarged localised reflections with dark discontinuities that divide them in two. This is remedied in the next step.

Marker extraction is performed on the source, target and sum images. A binary morphological dilation is performed on the sum image binary markers using a circular structuring element of diameter 5 pixels. This operation will enlarge the extracted markers, and also fuse sum-markers which have hairline discontinuities caused by a slightly longer time-interval.

The markers extracted from the sum-image can be used to determine which markers in the target image match the markers in the source image. For each marker in the source image, the marker in the target image that occupies the same sum-marker is identified as the matching marker, if the sum marker overlaps with these two markers exclusively.

Such a matching source and destination marker is used to determine the motion that the source marker underwent to reach the target marker position. The motion of the bubble is then approximated as the motion of its marker.

This algorithm is very fast and yields accurate per-bubble motion results. Figure 2.13 shows a sample of its output.

2.2.2 Region-based motion analysis

The per-bubble motion analysis calculates a motion vector for each bubble marker. Although per-bubble motion will prove valuable in motion-tracking of individual bubbles over extended periods of time, calculating statistics on this information proves problematic, as the motion information is extracted at irregular intervals. In addition, detailed mobility information is sometimes required for pre-determined sub-regions of the froth surface.

A Block Matching Algorithm (BMA) [21] was employed in order to calculate motion vectors at regular intervals in pre-determined regions of interest. This algorithm works by partitioning the source image into non-overlapping blocks and then attempting to match each of these blocks to

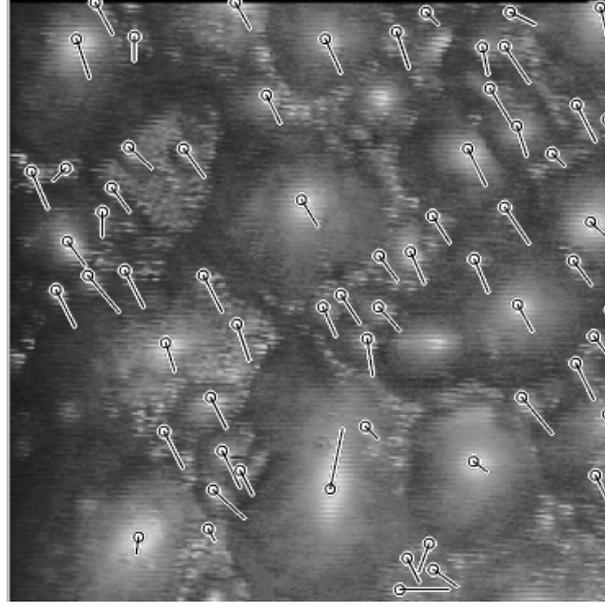


Figure 2.13: Sample results of per-bubble motion analysis

the corresponding block in the target image. Again, the source image and target images are of the same scene but taken a short time apart. The matching target block is found by calculating some similarity measure between the source block and several blocks in a target image search region. Of these candidate blocks, the one that achieves the highest similarity measure is selected as the matching block. The difference in position between the source block and the target block represents the motion that was undergone by this block in the time interval between images.

For this application, the Maximum Cross-Correlation (MCC) matching criterion was used. Each source block is cross-correlated with all the possible overlapping target blocks in its search region in the target image, and the position at which maximum cross correlation is achieved is selected as the position of the matching target block. The BMA with MCC matching criterion is a form of feature-detection by matched filtering [22, 23] where the matched filter in this case is the source block itself. The motion vector $\mathbf{V}_i(m, n)$ measured from image I_i to a following image I_{i+1} for a block centered at pixel location (m, n) is determined by:

$$\mathbf{V}_i(m, n) = (x_d - m, y_d - n)$$

where

$$(x_d, y_d) = \arg \max_{x, y \in \Omega} \left(\frac{\sum_{p=-M/2}^{M/2} \sum_{q=-N/2}^{N/2} I_i(m+p, n+q) I_{i+1}(x+p, y+q)}{\sqrt{\sum_{p=-M/2}^{M/2} \sum_{q=-N/2}^{N/2} I_i^2(m+p, n+q) \sum_{p=-M/2}^{M/2} \sum_{q=-N/2}^{N/2} I_{i+1}^2(x+p, y+q)}} \right)$$

The motion-block is M by N pixels and the search region Ω is a block congruent with the motion-block but larger. The normalisation is very important. It eliminates the effect of differing signal energies between the search block and the block in the search area that it's being matched with.

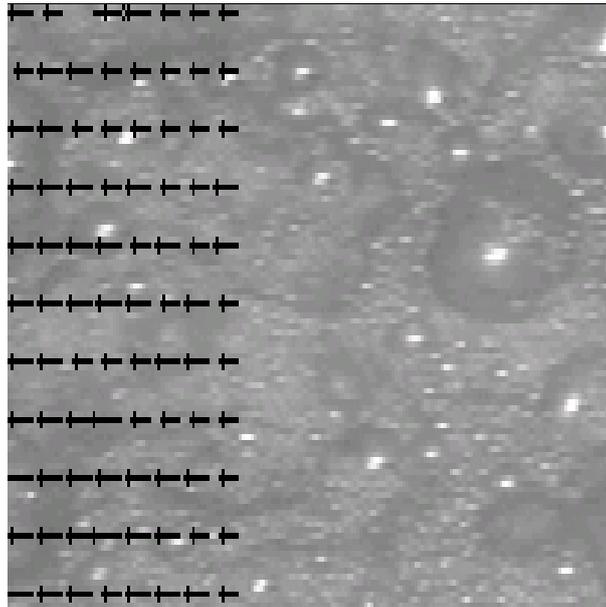


Figure 2.14: Output of the region-based motion estimation

This method is inaccurate if used on images with general motion, such as rotation, zoom and deformation [24]. In our case, the images are taken a very short time apart (20ms) and the motion blocks are kept relatively small. Under these circumstances and taking into account the nature of froth mobility we assume pure translational motion. Empirically, this turns out to be a valid assumption. Figure 2.14 shows sample output of the region-based motion-estimation algorithm. Each of the arrows represents an independently calculated motion vector.

The region-of-interest over which motion vectors are calculated, the size of the motion-blocks and the size of the search areas are all user-definable and should be calibrated to suit the physical setup of acquisition equipment. For our environment, the region of interest is set to cover a third of the image, the motion block size is 10x10 pixels and the search area is a 16x16 pixel congruent block. The motion block and search area sizes are constant for a region of interest. The effective image size is 380x142, which is either the odd or even field of a 380x284 image.

2.3 Conclusions

Table 2.1 shows average execution times of the main algorithm groups documented in this chapter. These timings exclude overheads such as image acquisition and display. The bubble segmentation includes a marker extraction, homotopic transform and watershed segmentation. The pixel analysis includes algorithms which return statistics on pixel intensities and colour information (these are documented in the next chapter). The algorithms were performed on 380x142 images

Analysis	Tank [s]	Scavenger [s]
Bubble segmentation	0.98	1.1
Pixel analysis	0.27	0.27
Bubble motion	1.03	1.56
Region-based motion	1.02	0.95

Table 2.1: Average execution times for various algorithm groups on tank cell and scavenger cell froth.

on a Celeron 366MHz processor, the region-based motion estimation region of interest was set to a third of the image, the motion block to 10x10 pixels and the search region to a 16x16 congruent block.

The tank cell froth consists of larger bubbles and the scavenger cell froth of many smaller bubbles. As to be expected, the bubble-segmentation and bubble-motion analyses take noticeably longer on the scavenger froth, whilst the pixel analysis and region-based motion analysis perform more or less the same on both types of froth.

The low-level image processing algorithms documented in this chapter extract valuable information that can be post-processed to yield higher-level parameters that more directly describe the flotation froth that is being examined. The following chapter will document this post-processing.

The fine-froth detection methods, the post-segmentation oversegmentation correction and the per-bubble motion algorithms are novel additions to flotation froth analysis. Although one of the fine-froth detection methods is not currently applicable, it indicates that there is much research potential in attempting froth-dependent segmentation of froth images.

Chapter 3

Parameter Extraction

In this chapter we document how the low-level image-processing information is processed and converted to usable froth surface parameters. This includes the calculation of discrete bubble-structure parameters, such as size and shape, for every distinct bubble, as well as the generation of summarising information, such as a histogram or a median value. In current automatic control systems, this summarising information is utilised rather than raw image-processing data. Table 3.1 lists all extractable parameters and available post-processing. Subsequent sections will discuss the available post-processing in detail.

3.1 Pre-processing of image-processing information

The per-bubble segmentation yields a tessellation of the froth surface image, where all the pixels belonging to a specific bubble are tagged as such in a two-dimensional array of the same dimensions as the image. In order to work with this data, it has to be converted to a bubble-ordered form (as opposed to a position-ordered form), i.e. where bubble segmentation information can be accessed on a per-bubble basis.

In order to do this, the tessellation is run-length encoded [22]. A list consisting of bubble-information atoms is formed. Each atom contains a list of line start positions and lengths which represents a rasterised bubble. This representation facilitates the calculation of the parameters detailed in the following sections.

Parameter	Statistics	Circular Statistics	Histogramming	Once-off
area/size	✓		✓	
blue	✓		✓	
bubble-count				✓
ellipticity	✓		✓	
fine froth ratio				✓
green	✓		✓	
intensity	✓		✓	
orientation		✓	✓	
perimeter	✓		✓	
red	✓		✓	
SNE				✓
valid froth ratio				✓
velocity (bubble)	✓	✓	✓	
velocity (area)	✓	✓	✓	

Table 3.1: Extractable parameters and available post-processing

3.2 Extractable parameters

Several parameters describing the physical properties of the froth surface can be calculated from the low-level image processing data. Most of these parameters are large vectors, describing e.g. either all the bubbles in an image or describing motion vectors in a region of interest, and will thus be further processed at a later stage by a combination of statistic calculation, circular statistic calculation and histogramming.

3.2.1 Area/Size

Operator-perceived bubble size is one of the most important parameters describing the froth surface. We approximate operator-perceived bubble-size with the bubble-area, i.e. the number of area units, or pixels, that the segmented two-dimensional projection of the bubble occupies. This is a safe assumption as absolute bubble-size is not required, but rather relative bubble-size. That said, the area of the two dimensional projection is very closely related to the true surface area of a bubble, with the estimation error relative to the true bubble size, the degree of curvature (more specifically the specific shape of the bubble-dome) and the height of the bubble if the projection takes place along the normal axis.

Bubble area is calculated by counting the exact number of pixels contained in its tessellated two-dimensional projection. This is done for each and every bubble in the image. The counting is done by performing a sum on the run length pixel information for that bubble. Converting the

area to more traditional units is a question of camera calibration. This is not required, as the relative measures can be used.

3.2.2 Bubble-Count

Bubble-count refers to the total number of bubbles present in a specific froth image. Areas that have been classified as fine froth by the fine froth detection algorithm in section 2.1.3.1 are not counted. The bubble-count can be used as an estimator of the average size of the bubbles in the froth image if the amount of fine-froth is also taken into account. If the amount of fine-froth is relatively small, the bubble-count is inversely proportional to the average bubble size.

More importantly, the bubble-count can be used as an indication of froth stability when continuous video is analysed. If the bubble-count shows dramatic variations in sequences of consecutive images, it indicates that in addition to the froth undergoing its normal motion, bubbles are bursting or getting created at a high rate, which indicates instability. If the bubble-count deviates only as much as is expected for normal motion, a stable froth is being examined. Bubble-count variations can be used as a stability indicator with longer image sampling periods but at considerably lower sensitivity.

3.2.3 Bubble-based Motion

The algorithm documented in section 2.2.1 results in a list of horizontal and vertical differential motion pairs for each marker or bubble that has moved. The horizontal and vertical differential motion pair is a two element vector describing how many pixels the marker has been translated in the time interval between the algorithm's two input images. Polar motion information, i.e. motion direction and motion magnitude, are calculated using these Cartesian vectors and are also stored as additional forms of the resultant motion. This list of per-bubble motion is post-processed at a later stage.

3.2.4 Colour/Intensity Information

The red, green and blue components of the full colour image are statistically collated and histogrammed at a later stage. Intensity information, which is the average of the red, green and blue components, is also calculated.

The intensity information is related to the type and concentration of the mineral being floated and is also related to froth grade [9]. For it to be usefully applied, it has to be calibrated. This calibration has to be continually updated, as the light conditions on the froth change continually.

Although it is calculated and available for use, no work was done on investigating or utilising its relationship with froth characteristics.

3.2.5 Ellipticity

For a two-dimensional shape S containing N binary-valued pixels, the centre of mass (\bar{m}, \bar{n}) is defined as [22]:

$$\bar{m} = \frac{1}{N} \sum_{(m,n) \in S} m, \quad \bar{n} = \frac{1}{N} \sum_{(m,n) \in S} n$$

and the (p,q) order central moments are:

$$\mu_{p,q} = \sum_{(m,n) \in S} (m - \bar{m})^p (n - \bar{n})^q$$

Utilising the run length encoded pixel data the centroids and the second order central moments, $\mu_{1,1}$, $\mu_{2,0}$ and $\mu_{0,2}$, are calculated. The covariance matrix for the shape (where each pixel is considered as a two-dimensional feature) can then be constructed as follows [25, p 138]:

$$\begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{bmatrix} = \begin{bmatrix} \mu_{20} & \mu_{11} \\ \mu_{11} & \mu_{02} \end{bmatrix}$$

Where σ_{11} and σ_{22} are the variances of the horizontal and vertical dimensions, and $\sigma_{12} = \sigma_{21}$ is the covariance (measurement of joint relationship) between the horizontal and vertical dimensions.

The best-fit ellipse for our two-dimensional shape, i.e. the binary blob-image of a bubble, is an ellipse with a second moment (covariance matrix) equal to that of the shape [22, p 394]. This is akin to approximating the probability distribution of the two-dimensional features, or pixels, with a bivariate normal distribution, in which case contours of constant density are ellipsoids centered at the centre of mass (\bar{m}, \bar{n}) with axes defined by $\pm c\sqrt{\lambda_i}\mathbf{e}_i$, where c is a constant defined by the constant density chosen, λ_i is the i 'th ordered eigenvalue and \mathbf{e}_i is its corresponding normalised eigenvector of the covariance matrix [26, pp 335–337].

We define the ellipticity ε of this two-dimensional ellipsoid to be the ratio of the longest axis to the shortest axis, i.e.

$$\varepsilon = \frac{|\pm c\sqrt{\lambda_a}\mathbf{e}_a|}{|\pm c\sqrt{\lambda_b}\mathbf{e}_b|} = \frac{\lambda_a}{\lambda_b}$$

In other words, eigenvalues are calculated for the covariance matrix of the segmented bubble-shape under examination, and the ratio of the largest eigenvalue to the smallest eigenvalue is calculated, as it represents the ellipticity of the best-fit ellipsoid.

This parameter is related to the viscosity and stability of the froth. An ellipticity of unity represents a circular bubble whilst a larger ellipticity represents a more elliptical or elongated bubble, which usually appears in a viscous stable froth [8].

3.2.6 Fine froth/Valid Froth Ratio

The total area of the regions classified as fine froth froth by the fine froth detection algorithm in section 2.1.3.1 is calculated as a ratio of the total image area. This is the fine froth ratio and can be used as a direct measure of the amount of fine froth in an image and consequently as an indication of bubble-size, especially when the ratio is large. When the ratio is relatively small, bubble-size information gained from segmentation is more significant.

The valid froth ratio is the ratio of valid froth area to total image area, where valid froth area is simply the difference in area between total area and fine froth area.

3.2.7 Orientation

The orientation is defined as the clockwise angle between the longest axis of the best-fit ellipsoid (see 3.2.5) and the horizontal. As discussed, the ellipsoid has axes $\pm c\sqrt{\lambda_i}\mathbf{e}_i$, with \mathbf{e}_i being the i 'th eigenvector of the shape's covariance matrix. The vector \mathbf{e}_m corresponding to the largest eigenvalue is used to calculate the angle, where the angle θ is defined as:

$$\theta = \tan^{-1} \left(\frac{\mathbf{e}_m[1]}{\mathbf{e}_m[0]} \right)$$

The angle is adjusted to be on $[0^\circ, 180^\circ]$.

The orientation can be used to determine to which degree ellipsoid bubbles are elongated in the direction of froth motion. The more the elongation is in the direction of motion, the more significantly it indicates viscosity.

3.2.8 Perimeter

The perimeter is the pixel perimeter of the segmented bubble-projection. A boundary tracing algorithm from [27] is performed on the bubble-segmented tessellation (all pixels belonging to a certain bubble are tagged uniquely, i.e. have a specific colour) of the froth surface image. The boundary tracing algorithm yields a list of pixel chain codes describing the Euclidean perimeter of the bubble.

Each chain code indicates the direction in which the next pixel can be found. There are eight directions, of which four are diagonal and four are vertical or horizontal. By traversing the list of chain codes and adding one pixel to the perimeter total for every horizontal or vertical chain code, and $\sqrt{2}$ pixels for every diagonal chain code, the Euclidean perimeter of the bubble can be calculated.

The bubble-perimeter is very closely related to the bubble-size.

3.2.9 Region-based Motion

The region-based motion estimation documented in section 2.2.2 results in a matrix of regularly spaced motion vectors all within the specified region of interest. The regular spacing is determined by the block size, as each vector denotes the motion that a specific partitioned block underwent from one image to the next.

These motion vectors are initially available as differential horizontal and vertical motion pairs, but polar representations are calculated as well. In contrast to the bubble-based motion estimation, summarising information (such as means, histograms, etc.) calculated from these vectors are more readily interpretable as they are regularly spaced.

Our parameter choices are documented in sections 2.2.2 and 5.2.

3.2.10 SNE

The SNE was discussed in section 2.1.2. This parameter is calculated once for a whole image. No post-processing is performed on it.

3.3 Statistics

All the multi-valued extracted parameters, i.e. parameters which have distinct values for each bubble, for each pixel or for each partitioned motion block in an image, are summarised with statistic values that are useful in online flotation control systems. These values are: minimum, maximum, standard deviation, mean and median. The mean refers to the arithmetic mean of the data and the median to an actual data value that has half the data values greater than it and the other half smaller than it. The standard-deviation is a measure of the spread of the data about its mean.

Before these values are calculated on the extracted data, parts of the data which might affect

the accuracy of the calculated statistics are discarded. Recall that all regions that are classified as fine froth are ignored in the calculation of any statistics or histograms. Furthermore, for the bubble area and perimeter parameters, the highest 0.5% and lowest 5% of parameter magnitudes are discarded when calculating statistic values or histograms. This is to eliminate any possible outlier data. Although the convention is to discard 5% from both the higher and lower ends, it was empirically found that 5% from the higher values discards important samples individuals when working with bubble area and perimeter. This is a form of the L-estimator defined in [28] where the L-estimate $\hat{\theta}$ is defined as:

$$\hat{\theta} = \sum_{i=1}^N a_i X_i$$

with X_i being the sorted series of values to be estimated and a_i a corresponding set of coefficients. We define a number of the first and last a_i coefficients to be zero, and the remaining ones to be $\frac{1}{N-d}$ where d is the total number of a_i coefficients that are set to zero. In this way, we eliminate outliers before calculating our estimates.

With smaller bubbles where we find that less than 10 pixels can constitute a distinct segmented bubble, the best-fit ellipsoid that is used to calculate bubble orientation and ellipticity tends toward circularity. In most cases, this is an inaccurate approximation that takes place because of spatial under-sampling, and we do not want these results to bias our statistics or histograms. For that reason, we discard the orientations and ellipticities belonging to the 30% smallest bubbles. In addition, 0.5% of the largest bubble ellipticities are discarded to account for outliers.

With colour and intensity data, it was found that eliminating 2.5% from the greatest and the smallest magnitude data points compensated for the localised reflections found on the bubbles and for the darkest outliers. Intensity and colour information should not be affected by the localised reflections, as these are artifacts of the light sources. Figure 3.1 shows a froth image and then an image consisting of just the top 2.5% of its intensity values. Note that these include many of the brightest structures, but that 2.5% is a very modest estimate. Notice in the histogram the clustering at the highest and lowest values, representing the values that we discard.

None of the other parameter sets are culled in any way. Discarding angular data such as bubble orientation is not easily justifiable, as the definition of an outlier is absolutely dependent on where the reference diameter is chosen. Table 3.2 summarises the pre-processing that takes place.

The calculation of these statistic values is different for linear data (such as bubble area) and for angular data (such as bubble orientation). Table 3.1 shows which methods were used for which parameters. A mark in the ‘‘Statistics’’ column indicates that normal linear statistics are calculated, and a mark in the ‘‘Circular Statistics’’ column indicates that special circular calculations were performed. The next two sections outline the specific calculations that were used.

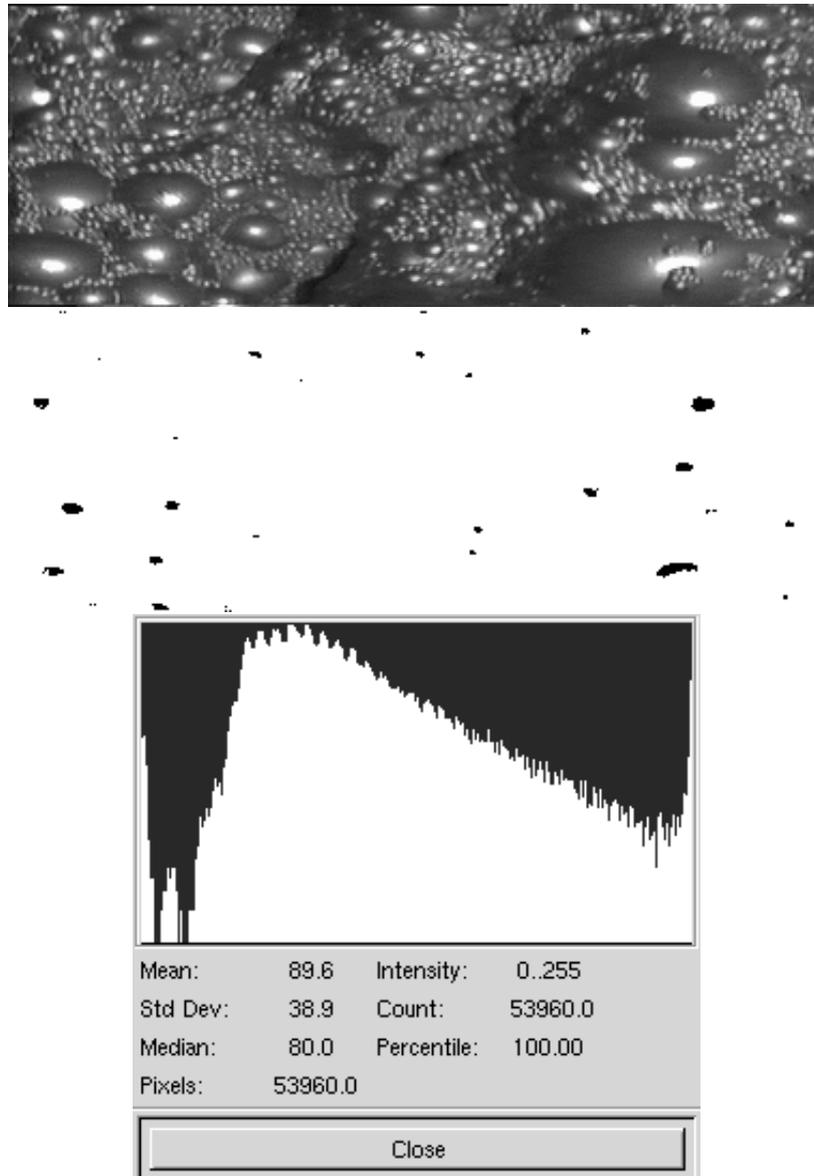


Figure 3.1: Froth image, image of top 2.5% of intensity values and histogram of froth image.

Parameter	Fine Froth Areas	Culling (low %/high %)	Extra culling
area/size	discarded	0.5/5.0	
intensity/colour	N/A	2.5/2.5	
ellipticity	discarded	0.5/0.0	30% smallest bubbles
orientation	discarded	0.0/0.0	30% smallest bubbles
perimeter	discarded	0.5/5.0	
velocity (bubble)	N/A	None	
velocity (area)	N/A	None	

Table 3.2: Pre-processing of parameter arrays

3.3.1 Linear Statistics

The array of data points is sorted in ascending order of magnitude and any necessary culling is performed (this is an L-estimation, as discussed in section 3.3). The first element in the sorted and culled array is stored as the minimum, the last element as the maximum and the middle element as the median. The arithmetic mean is calculated by summing all the data points and dividing by the total number of elements. The sampled standard deviation σ is defined as

$$\sigma = \frac{1}{N-1} \sqrt{\sum_{i=0}^{N-1} (x_i - \mu)^2}$$

where x_i is the i 'th data element, N is the number of elements and μ is the statistic mean. In this case, we used the median as a robust estimate of the statistic mean.

3.3.2 Circular Statistics

When considering the example of two angle samples at 359° and 1° , one notes that the arithmetic average is 180° and the arithmetic deviation is also way too large. One would expect the “true” mean to be 0° , and the variance more or less 1° about the mean.

If we selected the y-axis as being the 0-direction, the arithmetic answers would be more accurate, but this dependence of linear measures on the choice of zero direction indicate that they are not appropriate for circular distributions.

For the circular statistics, i.e. bubble orientation, bubble-based motion angle and region-based motion angle, special care has to be taken. It was decided that the minimum and maximum values would simply be the numerical minimum and maximum angles. The calculation of the circular mean, median and standard deviation is discussed in appendix A. Although the methods

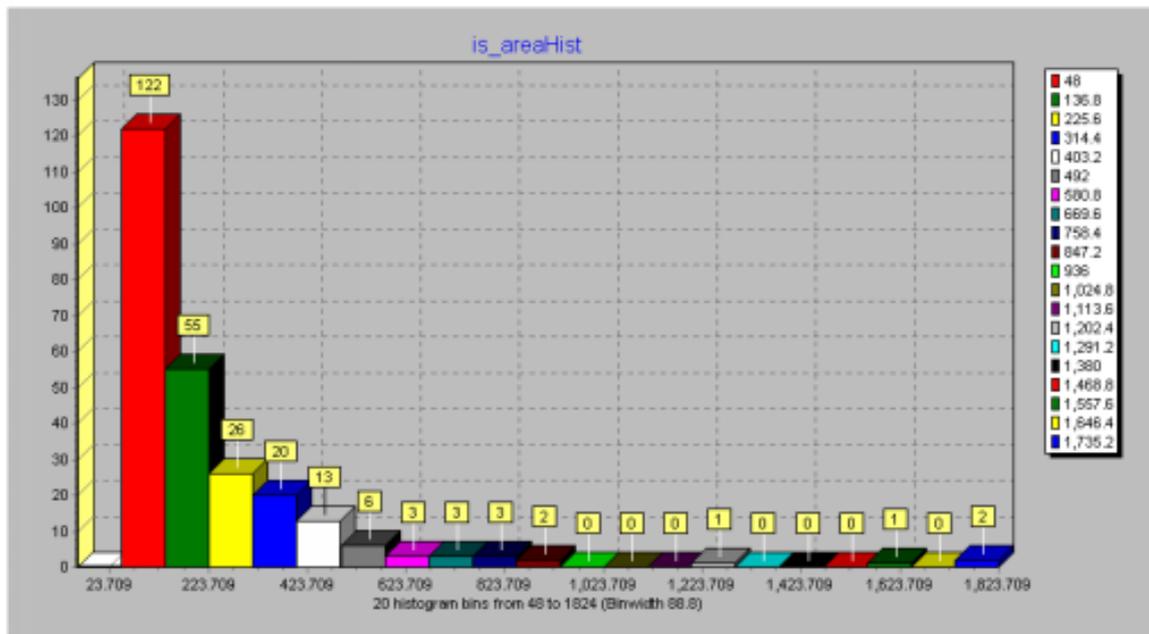


Figure 3.2: Example of area histogram

of calculation differ from the linear case, the meanings of the statistic values remain conceptually the same.

3.4 Histograms

Whereas the statistic calculations conveniently yield a number of scalar summarising parameters for each image that are easily interpretable, it is sometimes necessary to work with a more detailed summary, such as a non-accumulative distribution, commonly known as a histogram. These can be used for classifying a type of froth, akin to the work done by Moolman et. al [11,8] and the JKMRC [16] as documented in section 1.2. In addition, an operator monitoring the flotation performance can more readily make detailed deductions about the nature of the froth from a bubble size histogram than from the bubble size mean or standard deviation.

The histogram is a plot of the number of occurrences vs. a bin number. Each bin number represents a certain range of values for the data being examined. One can thus make observations about e.g. which percentages of the bubbles present are medium-sized or very large and so forth. The parameters for the histogram calculation, namely the number of equal-size bins, the start value and the end value are all user-definable. Figure 3.2 shows an example of an area/size histogram.

3.4.1 Histogram Accumulation

Histograms can be calculated for each image that is captured and analysed. These analyses usually take place at regular intervals, and with each new analysis, the histogram of the previous analysis is replaced with a new one. Instead of this behaviour, histograms can be accumulated over time to yield an accumulative histogram which represents the nature of the froth over a certain time interval. In this way, the accumulated histogram is able to contain information about past perturbations in a flexible way by acting as an estimator and in fact smoothing histograms.

The accumulation is performed with an exponential forget factor. Let \bar{h}_n and h_n be the accumulative and instantaneous histograms at discrete time n and λ an adjustable forget-factor that satisfies $|\lambda| \leq 1$, then:

$$\bar{h}_n = \sum_{k=0}^n \lambda^k h_{n-k}$$

This can be implemented as a recursive system as follows:

$$\bar{h}_n = h_n + \lambda \bar{h}_{n-1}$$

Note that a forget-factor $\lambda = 1$ means that no forgetting takes place, i.e. all past histograms contribute equally to the current accumulative. The smaller the forget-factor, the more the accumulation discounts the effect of past histograms, up to the limiting case where $\lambda = 0$ and only the current histogram makes a contribution.

Such an accumulative histogram can be used to examine the non-instantaneous nature of the froth. The forget-factor ensures that effects of process disturbances can be gradually discounted in the accumulative histogram. When the accumulated histogram is examined, it is useful to scale it as follows:

$$\frac{1}{1 + \lambda + \lambda^2 + \lambda^3 \dots + \lambda^n} \bar{h}_n = \frac{1 - \lambda}{1 - \lambda^{n+1}} \bar{h}_n$$

This yields a histogram which represents the accumulative nature of the froth, but whose values are in range for a normal instantaneous histogram.

Chapter 4

Practical Implementation

In this chapter the practical implementations and use of the computer vision froth surface analysis algorithms are detailed. In order to document the context in which these implementations are applied, a “generic” froth flotation based extraction plant is described, as well as the complete flotation control system of which one of the surface analysis implementations forms an integral part.

4.1 Froth Flotation Plant Basics

The setup of a flotation plant is naturally dependent on the type of mineral that is to be extracted. For any one mineral there are also several configurations that can be employed. Glembotskii [2] calls these configurations schemes, and defines them as combined systems of different flotation operations, where the scheme used depends on the properties of the mineral that is to be extracted, the desired quality of the product and naturally on economic factors. Such a scheme can consist of one or more stages, where a stage is defined to be a crushing operation and the flotation operations following it.

During the crushing operation ore is crushed or milled in order to liberate the valuable particles. In froth flotation plants, the subsequent flotation operations (as described in section 1.1.1) are implemented using flotation cells. These flotation cells are large containers with devices for injecting air-bubbles into the slurry and promoting motion in the pulp in order to facilitate air-bubble mineral-particle collisions. Mineral-laden bubbles float to the top and form the froth that can be skimmed off and filtered to form the valuable mineral concentrate.

This concentrate can go through any number of additional grinding and flotation operations (flotation cells) before it becomes the final concentrate. The tailings, i.e. particles which are

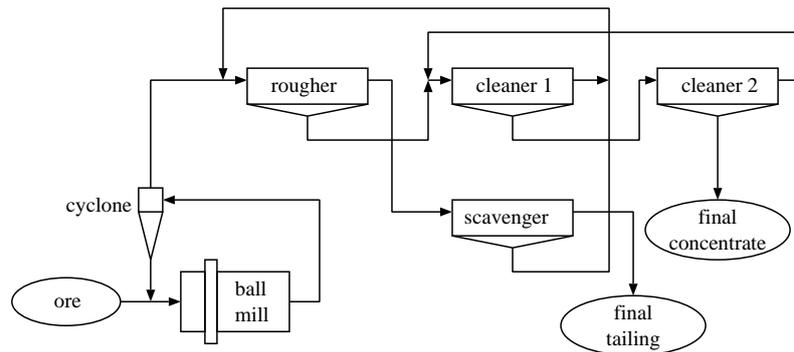


Figure 4.1: Diagram of simple froth flotation circuit

not removed in the concentrate in a flotation cell [3], can also be reground and/or passed through additional flotation cells. Eventually, final concentrate and final tailings outputs are attained.

Figure 4.1 is a diagram of a very simple froth flotation circuit from [3]. The rougher, scavenger and cleaners are flotation cells and their names are assigned according to their purpose and position in the circuit. Note that output from the bottom of a cell is concentrate and output from the right indicates tailing. The cyclone is a device that separates particles on the basis of fineness. Here it is used to return particles that aren't fine enough to the ball mill.

4.2 Fluxar Flotation Control System

Industrial colour cameras are mounted at various points in the flotation plant in positions where they can monitor the surface of the flotation froth in flotation cells. Custom-manufactured metal camera enclosures are pyramidal and reach down to near the froth surface as illustrated in figure 4.2. Their purpose is to eliminate ambient light interference as much as possible, thus increasing the accuracy of the image processing algorithms. In addition, a strong spotlight is mounted next to the camera. Figure 4.2 shows an example of the Fluxar camera setup at a South African platinum extraction plant. Whilst systems like the one by the JKMRC [16] claim to work with natural light, they also make use of shade flaps and spot-lights to increase accuracy.

The cameras are of the interlaced kind, i.e. the images they return consist of two interlaced images¹ taken 20ms apart (in the case of PAL). The even and odd images that constitute one interlaced image are extracted and used as separate images for dynamic analysis.

An obvious yet very important issue is that of acquisition hardware quality and any other measures taken to improve the quality of the acquired images. The importance of utilising the highest

¹All the even lines belong to one image and the odd lines to the other.

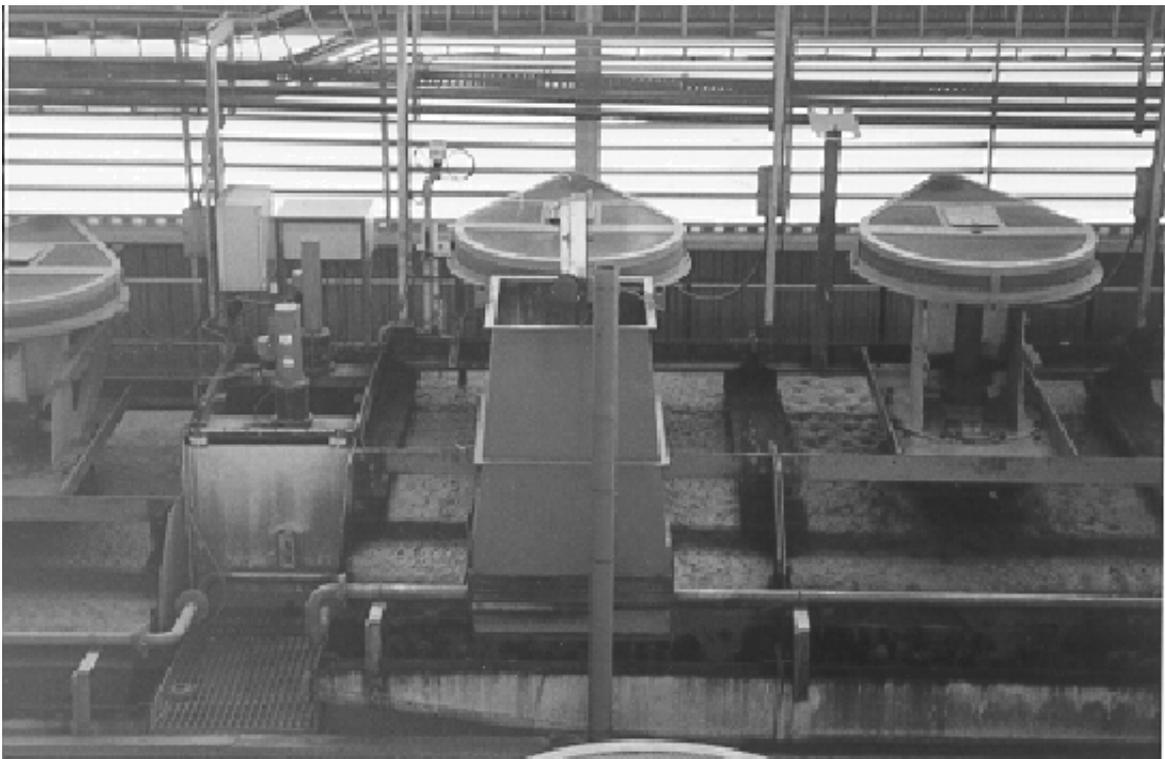


Figure 4.2: Camera setup

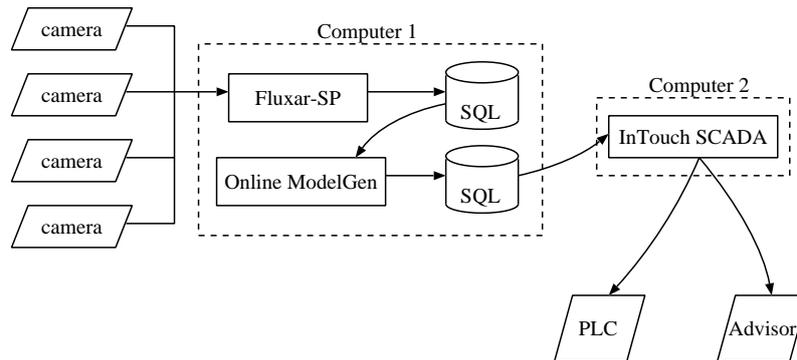


Figure 4.3: Diagram of Fluxar Flotation Control System

quality cameras, cabling and frame grabber cards can not be overstated. The price of these items is small in comparison to the gains brought about by increased parameter accuracy. All possible steps that can improve the lighting conditions and camera stability should be taken. The image analysis algorithms are relatively robust, but improved input equates to more accurate output.

PAL video cables run from the cameras to the control room where the computer hardware is situated. These video signals are sampled by Fluxar-SP, the full-featured froth analysis implementation discussed in section 4.3, which performs any combination of configurable froth analyses on this video input. At the moment, dynamic analysis is performed on the even and odd fields of a single image. These fields are exposed 20ms apart, as per the PAL standard. It stores all the froth information that it extracts into an SQL database, from whence it is accessed by Online ModelGen, an application which performs filtering, basic mathematical manipulation and fuzzy model decision making on this information. This output is stored in another SQL database, where it is accessed by Intouch SCADA (Simultaneous Control And Data Acquisition) and used to perform automatic control via the plant's PLCs and/or to advise the plant's operators about recommended control actions. Figure 4.3 shows a block diagram of the Fluxar Control System. The reader is referred to the Crusader [29] and Fluxar [30] web sites for further information.

4.3 Implementations of Froth Surface Analysis

The image analysis of chapter 2 and the parameter extraction of chapter 3 were implemented in several distinct applications with differing functionality and applicability. Most of the analysis functionality was abstracted in a shared analysis kernel, which will be discussed first.

4.3.1 Analysis kernel

The analysis kernel is a body of program code that implements most of the froth surface analysis logic in a platform-independent and reusable manner. It consists of object oriented C++ source code and is primarily developed on the Linux operating system, due to its standards compliant development tools and overall stability. Note that the analysis kernel source code remains platform-independent, i.e. it can be recompiled on other operating systems and processor types without modification.

4.3.2 Fluxar-SP

Fluxar-SP is the Fluxar Sensor Platform, a general-purpose froth surface video acquisition and analysis application. It is the most full-featured application of the froth analysis kernel and offers its operator access to all the image analysis and parameter extraction algorithms shown in table 3.1. It was designed primarily for deployment as the measurement component of the complete Fluxar flotation control system, which is documented in section 4.2, but is also very valuable as a research tool.

The package is able to capture images from live video input from up to four channels and perform any combination of analyses and parameter extractions on the captured images. The extracted information is stored in a remote or local SQL database for use by the other components in the flotation control system. All parameters of the acquisition and analysis are configurable by the operator through the user-friendly graphical user-interface. Figure 4.4 shows some of Fluxar-SP's user-interface elements.

Operators are able to view live video from the current acquisition channel and examine calculated histograms and extracted data as they are generated. Figure 3.2 shows a sample histogram as displayed by the package. Several output images are also shown as they are generated, namely an image with bubble segmentation lines overlaid, an image with region-based motion vectors overlaid, and an image with bubble-based vectors overlaid. A screen-shot of a sample image display window is shown in figure 4.5. This visual feedback enables the operator to quantitatively judge the performance of the algorithms. The operator can also view the values of scalar parameters as they are extracted by making use of the database viewer shown in figure 4.6.

As mentioned earlier, Fluxar-SP offers a very convenient way of configuring the various available algorithms, thus facilitating its operation by plant operators. In addition, this functionality makes it a valuable tool for experimentation with the froth analysis algorithms. Figure 4.7 shows the analysis options dialogue box. Here the user is able to adjust the motion estimation parameters by selecting the region of interest with a mouse and by entering block size and search area parameters using the supplied controls. The marker bubble area ratio threshold (see section 2.1.3.1) is adjustable and the oversegmentation correction (see section 2.1.1.4) can be activated or

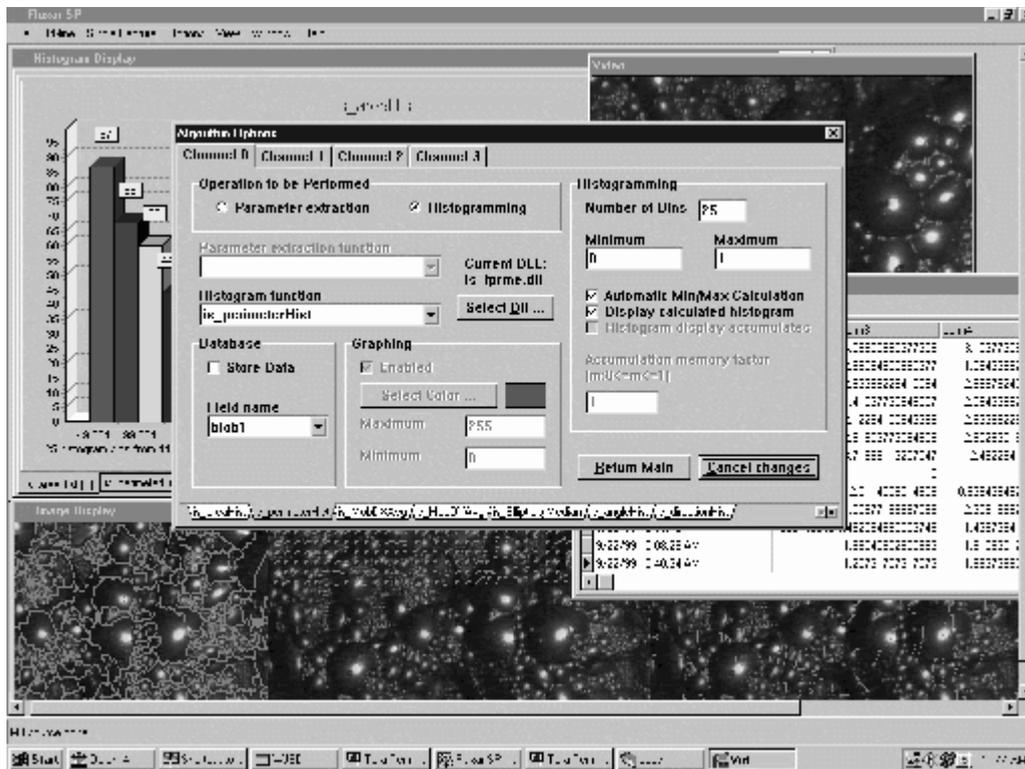


Figure 4.4: Screen-shot of Fluxar-SP

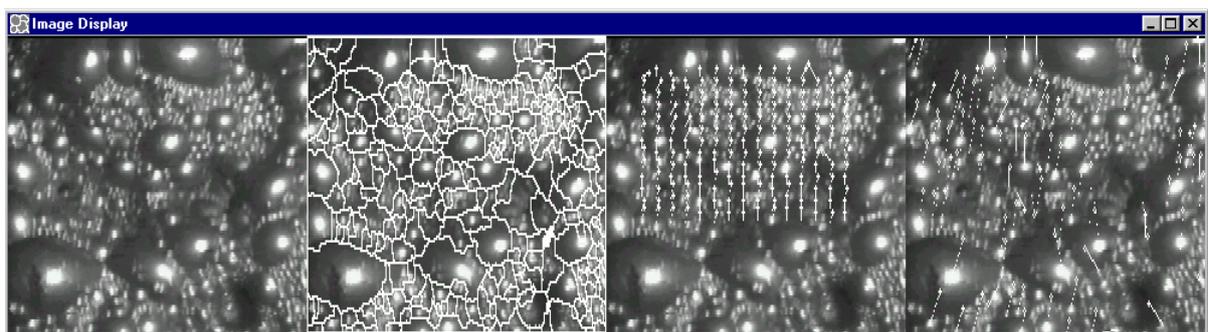


Figure 4.5: Screen-shot of sample Fluxar-SP image display window. From left to right: original, bubble segmentation, region-based motion, bubble-based motion.

Source 0	Source 1	Source 2	Source 3		
DateTime	NumBubs	AreaSD	AreaAvr	AreaMed	EllipSD
6/18/99 8:30:00 PM	27	6160058113	415.6	300	3850423409
6/18/99 8:30:11 PM	4	5727331551	24348.5	25536	50809.42229
6/18/99 8:30:21 PM	41	7453810838	2162162162	104	6962259045
6/18/99 8:30:30 PM	5	1522501505	447.6	244	1829541051
6/18/99 8:30:41 PM	11	7055903281	10909090909	280	5570526024
6/18/99 8:30:50 PM	9	4506943616	4444444444	272	5922844743
6/18/99 8:31:01 PM	10	7762988467	158	174	8769407165
6/18/99 8:31:10 PM	38	9175818877	4444444444	118	8762179575
6/18/99 8:31:21 PM	10	8928051153	6547.6	460	3323.774199
6/18/99 8:31:31 PM	6	2285222023	308	278	8208799109
6/18/99 8:31:41 PM	19	0322272532	8947368421	228	2697271003
6/18/99 8:31:51 PM	22	1545159942	251.1	190	5710085817

Figure 4.6: Screen-shot of Fluxar-SP database viewer with sample parameters

deactivated. Note that these parameters can be different for each acquisition channel. Figure 4.8 shows the algorithm options dialogue, where the user may configure combinations of parameters and/or histograms to be calculated for each acquisition channel.

It is able to acquire images from up to four multiplexed video-channels using either a Matrox Meteor or IntegralTech FlashPoint 3D Plus frame grabber card. As an added feature it is able to acquire from bitmap images stored on disc. As the on-line acquisition methods can be configured to store images to disc as they capture, these images can later be analysed off-line, in order to diagnose process disturbances after the fact. Note that in addition to being a crucial part of the automatic flotation control system, Fluxar-SP is a valuable research tool.

The Fluxar-SP front-end was developed using Borland Delphi 4.0 and this instance of the analysis kernel built with Borland C++ Builder 3.0. Development started on an initial third-party attempt at a flotation froth sensor-platform. More than 75% of the functionality was developed as part of this work.

Fluxar-SP is used actively in research environments in South Africa and Australia and has been efficiently functioning as the sensor component of a flotation control system in advisory mode on a real platinum extraction plant in South Africa since the end of 1998, and as the sensor component of a closed-loop flotation control system on this plant since the middle of 1999.

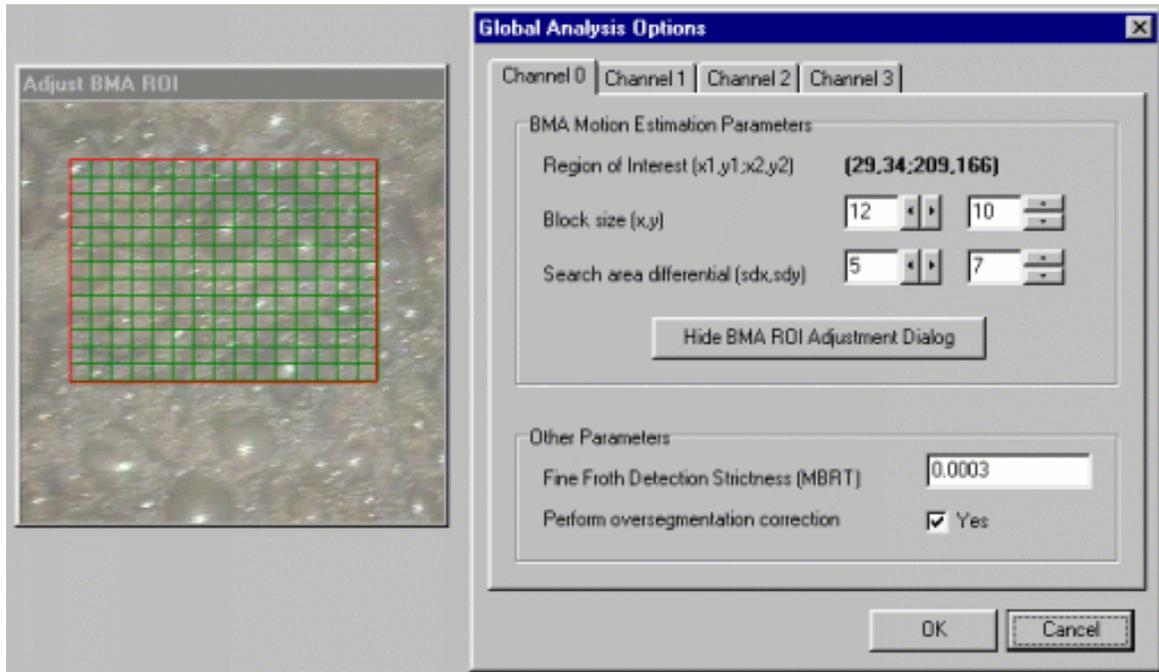


Figure 4.7: Screen-shot of Fluxar-SP Global Analysis Options dialogue

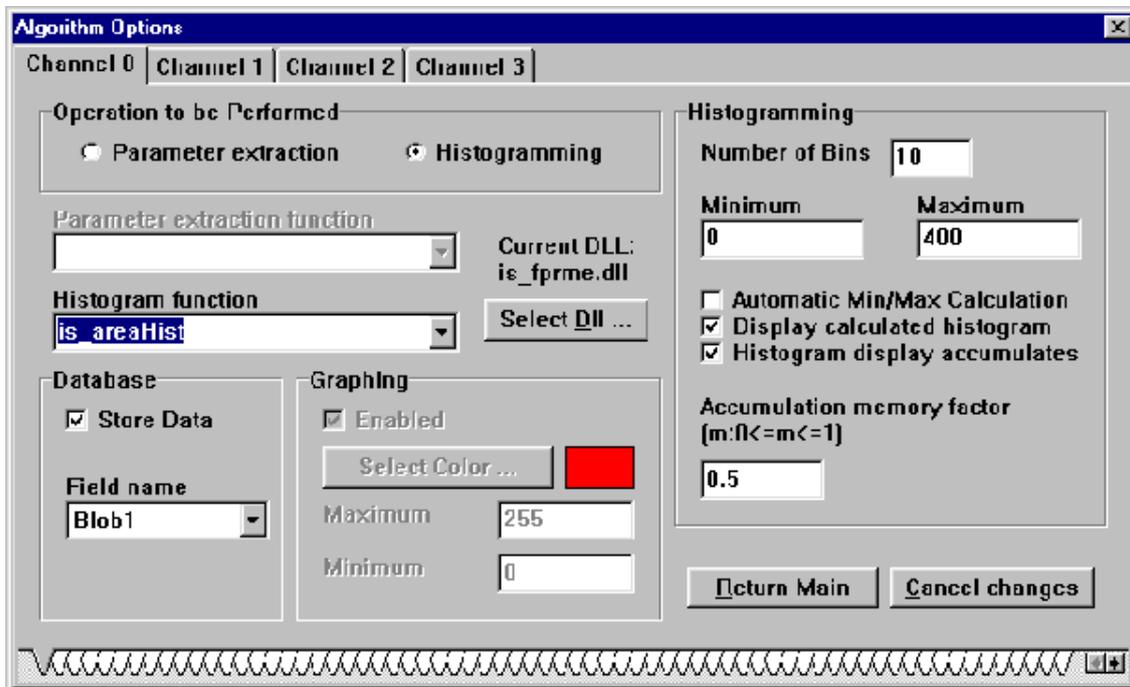


Figure 4.8: Screen-shot of Fluxar-SP Algorithm Options dialogue

4.3.3 Fluxar-Speed

Froth stability and mobility are very important measurements in automatic flotation control [9]. This importance was realised in another application of the froth analysis kernel. FluxarSpeed is an embedded application that is able to perform region-based motion analysis in real-time and output average motion as an analogue signal, as values from its serial port and as a continuous reading on its mounted current meter. Depending on the configured motion analysis parameters, the prototype system can analyse up to 15 images per second.

It was designed for installation in flotation plants which do not require the extended functionality of Fluxar-SP. A flotation plant would mount several of these units at different positions and use the analogue outputs as measurement signals in their existing control platforms. By keeping froth mobility at predefined “good” levels, extraction plants can be run at greater efficiency.

FluxarSpeed consists of x86 architecture hardware running the Linux operating system. Figure 4.9 shows the internal layout of the FluxarSpeed prototype. At bottom left the Intel Pentium III 450MHz CPU can be seen, and at top left the built-in greyscale camera and PCI grabber card.

When installed, the system can be configured via its built-in network interface by using telnet, or via one of the serial ports, by connecting a portable computer with a serial terminal application. All the parameters of the area-based motion analysis can be configured in this way. The software on the product can also be updated using these methods. The system is designed for full-time low-maintenance operation, whilst maintaining a very high-level of in-place configurability and upgrade-possibilities. This makes for a low-cost high-gain measurement device.

Outokumpu Mineral Processing Oy of Finland will be marketing the Fluxar-Speed system on an international basis. For this purpose, the system is being ported to an industrial embedded PC platform. The complete software aspect of FluxarSpeed was engineered as part of this work.

4.3.4 Continuous sampling platform

The third implementation of the froth analysis kernel was developed purely for research purposes. It consists of a continuous sampling component and a batch analysis component. The continuous sampling component is able to capture video sequences as continuous sets of distinct images and store them on disc. The batch analysis component is able to perform the full froth static and dynamic analysis on these stored image sequences in an off-line fashion and store the extracted parameters on disc.

Nowhere in documentation has this type of continuous froth video analysis been performed. It is especially important in examining parameter trends which occur at a higher frequency than the customary sampling rate in existing research and industrial setups. This software is able to

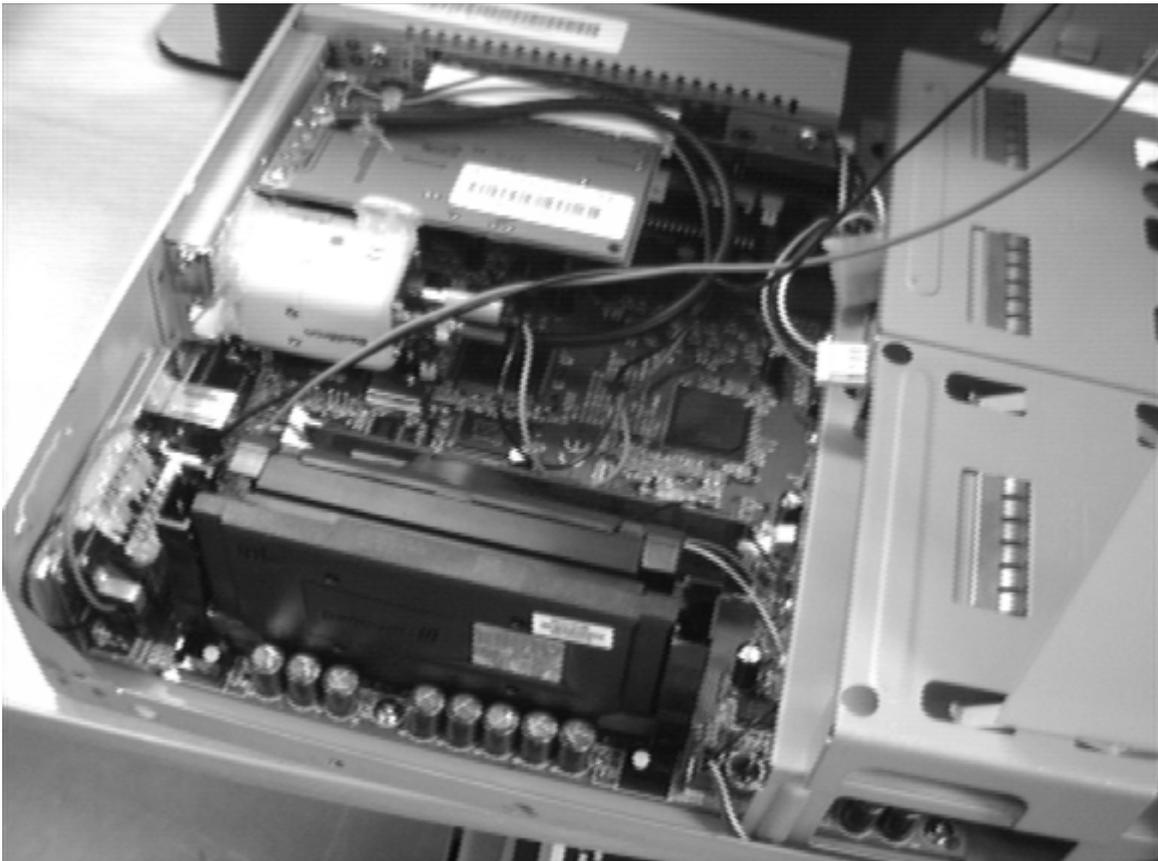


Figure 4.9: Photo showing FluxarSpeed prototype internal layout

sample at periods that are an integral multiple of 20ms, whereas the customary sampling rate for industrial setups is 30 to 90 seconds. The results of this novel experiment are detailed in chapter [5](#).

Chapter 5

Performance and Impact

This chapter details the performance of the flotation froth analysis algorithms and the impact of their use on real flotation plants. The difficulty of *conclusively* validating algorithm performance is explained and investigations into alternative measures of accuracy are documented. Experiments with continuous video flotation analysis and their pertinence to algorithm validation are documented. Subsequently, the performance of the algorithms in real-world flotation-plants is discussed. Naturally, this practical measure is a very good indication that the underlying algorithms are performing well in their function as measurement component of the control system.

5.1 On Algorithm Validation

Quantifying the image processing accuracy of the flotation algorithms is a very difficult task, as there are no benchmark algorithms to compare with nor existing databases of froth images and “correct” image parameters. Creating such a database is not in the scope of this work, although it would be a very valuable asset in the light of current developments in flotation froth analysis. An alternative is to compare the algorithm performance to that of a human vision system. This is also not viable, as visually inspecting a statistically significant sampling of images and comparing the result of this inspection with the algorithms’ performance on the sampling would take a prohibitively long time.

5.1.1 Experiments with synthesised motion

Another method of testing computational vision algorithms is synthesising data with known parameters and then testing the algorithm extracted parameters against the known ones. Once again

Algorithm	Horizontal MSE	Vertical MSE	Distance MSE
Bubble Motion	1.2298	2.9632	3.1416
Area Motion	0	0	0

Table 5.1: Mean square pixel errors for horizontal, vertical and Euclidean motion calculated over 100 test images.

this is a difficult undertaking with froth surface images, as the motion that they undergo is complex and varied. Also, synthesising images showing the same structural properties that static froth images do would be an impossible task.

However, in order to test our particular implementation of the block matching algorithm and to validate partly the bubble-based motion estimation method, random pure translational froth motion was synthesised. An image is created which is four times the size (double the width and double the height) of an existing froth surface image. The top left quadrant of the large image is a copy of the froth surface image, the bottom left is a vertical mirror image of the froth surface image, and the right two quadrants are a horizontal mirrors of the left two quadrants. This image can be used to generate arbitrary sequences of desired translational motion by “rolling” it vertically and or horizontally. If these sequences were generated by taking sub-images of an unmodified froth image, the random sequence would be constrained by the dimensions of the image. Both the region-based and bubble-based motion estimation implementations were tested on such sequences, each consisting of 100 images with random horizontal and vertical motion throughout the sequence and the magnitude of the motion equal or less than $\sqrt{50}$ pixels.

As to be expected, the region-based BMA motion estimation generated perfect results. The experiment only serves to demonstrate the correctness of our implementation. In the case of the bubble-based motion estimation however, the experiment served as an effective validation and rating of its performance.

Table 5.1 shows the mean square pixel errors for horizontal, vertical and Euclidean motion as calculated by the bubble-based motion and area-based motion estimators. MSE is defined as

$$\text{MSE} = \frac{1}{N} \sum_{n=0}^{N-1} |\mathbf{x}[n] - \mathbf{x}_e[n]|^2$$

where N is the number of values, \mathbf{x} is a vector containing the correct motion values and \mathbf{x}_e is a vector containing the algorithm estimated values. Recall that the bubble-based motion estimation attempts to return a motion vector for every bubble, i.e. at irregular intervals. In spite of this, the numerical average gives a very good indication of the global motion that has taken place.

This experiment also shows up a weakness in the algorithm which can be attributed to the size of the structuring element that it used for the binary dilation of the summed marker image (see section 2.2.1). It was found that the largest errors occur at samples where the horizontal or vertical component is 4 pixels or more. If the structuring element were to be made larger to

0 motion of a greater magnitude, more incorrect bubble-motion vectors would be generated in regions with smaller bubbles.

The maximum detectable magnitude of motion detectable by the region-based motion estimation is limited by the size of the search area. This size is adjustable however. Although making it larger increases accuracy, the size increase increases execution time significantly.

5.1.2 Qualitative judgement

Extensive qualitative measurements of the algorithms' performance were made by examining its graphical outputs both on the extraction plant and in the laboratory with hours of recorded flotation froth video footage. The graphical outputs are segmented images such as that in figure 2.8 and images with calculated motion vectors overlaid such as that in figures 2.13 and 2.14. Judging by these observations, the algorithms perform satisfactorily.

5.2 Continuous Sampling Experiment

When the froth surface analysis software is in use on the plant, typical image sampling periods are at the very least 15 seconds, more often 30 to 120 seconds. This is also the norm for JKFrothCam [16]. This period is very short in comparison to the time constants of processes on the plant where a disturbance is considered very fast if it lasts an half-hour. Because of this, no work has been published on the analysis of continuous sampling and analysis of froth surface video data. In the course of this work, some investigations were made into this.

Video from a laboratory flotation cell and from four industrial cells in a South African platinum extraction plant was used. Segments of 16 seconds (representing 400 images at 25 images per second) were digitised, as well as longer segments of 2 minutes (representing 1500 images at 12.5 images per second). Full static and dynamic analyses were performed on some 19200 images. The continuous sampling platform of section 4.3.4 was used to perform the digitising and analysis. Interlaced images of dimensions 380x284 pixels were captured, of which only the even fields were used for these analyses. The region of interest for the area-based motion analysis was chosen on the grounds of the location of the weir which varies for our datasets, but was kept to a third of the image area. The effective partitioning (motion block) size for the region-based motion estimation was chosen as 10x10 pixels and the search area as a 16x16 congruent block. On an entry-level Intel Celeron 366MHz computer, a full static analysis on one image and a full dynamic analysis on that image and the next takes on average 2.5 seconds.

On the grounds of the investigation of the analyses results, several important observations can be made.

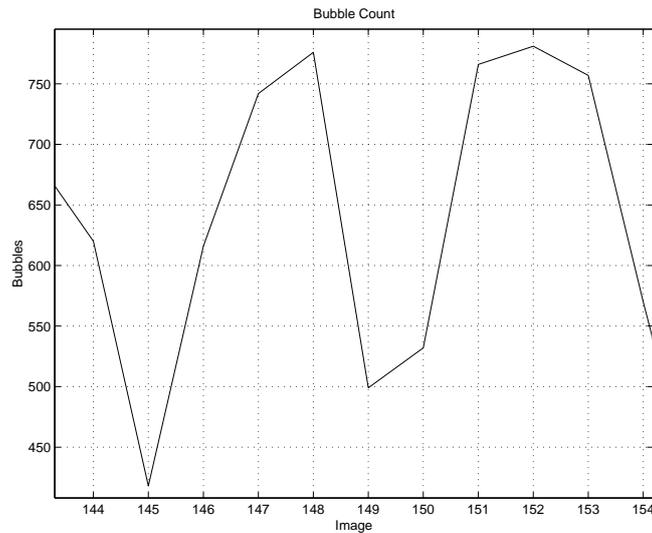


Figure 5.1: Small section of number of bubbles plot for 25 image per second sample set of laboratory flotation cell froth illustrating number of bubbles sensitivity.

5.2.1 Continuous Video Static Analysis

In this section, we will look at the parameters extracted from images taken of flotation froth in a laboratory environment. This froth is quite unstable and facilitates the investigation of selected froth dynamics and events which aren't so readily available in controlled froths in industrial environments.

The bubble count and average bubble size (by its inverse relation to number of bubbles) can show relatively large jumps for very little humanly visible change between two images. If it is not immediately apparent why a parameter experiences a discontinuity, examination of other parameters often yields the answer, and this will shortly be demonstrated.

For example, figure 5.1 shows a small section of the number of bubbles parameter plot for a 25 image per second sample set of the laboratory flotation cell froth and figure 5.2 shows some of the segmented images corresponding to the graph.

Image 144 shows an image containing a few very large bubbles and a majority of very small bubbles. Figure 5.3 then shows a sharp acceleration in vertical motion which is reflected in the ellipticity plot (the average ellipticity increases because of the froth's natural cohesive resistance to the acceleration) in figure 5.4 and in the froth image 145. Probably because of this sudden perturbation, the large bubble on the right starts to weaken and is drained of its mineral load. Note how the segmentation reacts to the start of the bubble-burst. In image 146 this large bubble bursts and the bubble count increases dramatically with a corresponding decrease in ellipticity (bubble tension is decreased by the large bubble burst) and in vertical speed. A brief steady state

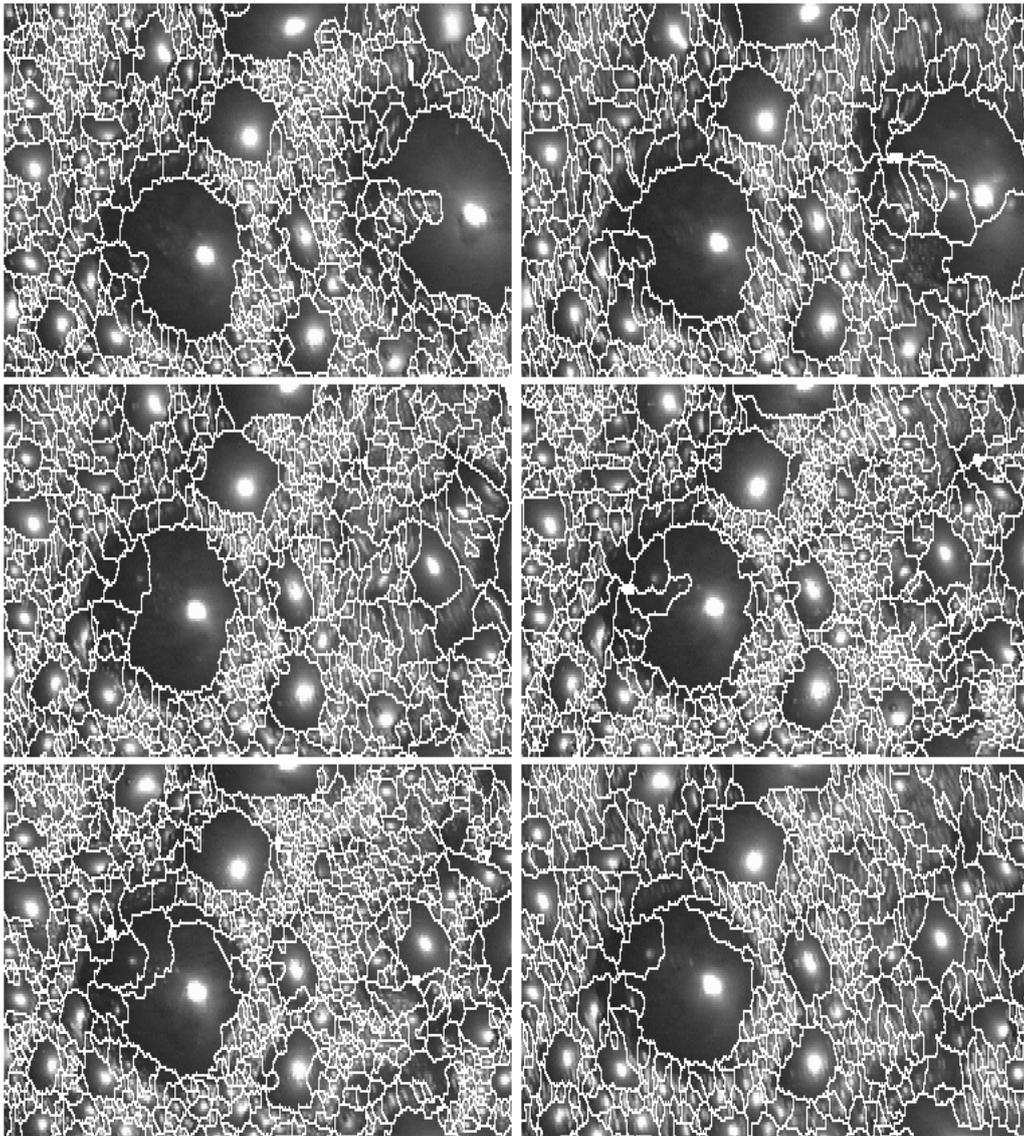


Figure 5.2: Segmented froth images corresponding to graph in figure 5.1. Images 144–149 left to right, top to bottom.

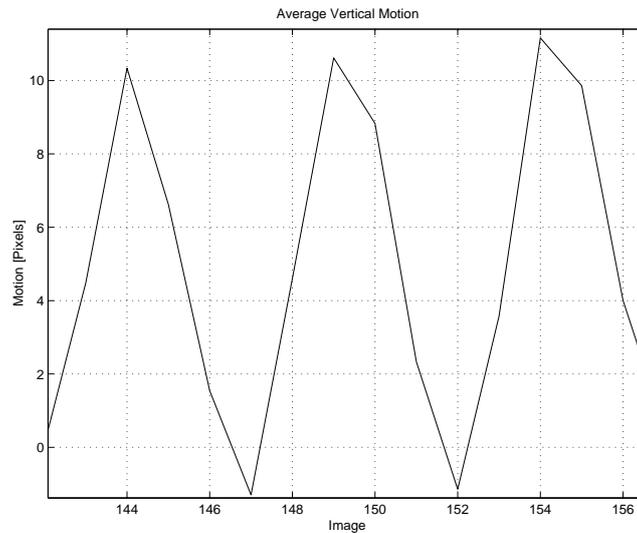


Figure 5.3: Section of average vertical pixel motion as calculated by region-based motion calculation for 25 image per second sample set of laboratory froth. Compare with the images in figure 5.2.

is reached at image 148 as the froth almost comes to a standstill and ellipticity falls back to its initial value. By image 149, the froth accelerates once again and this is reflected in the increased ellipticity, decreased bubble-count and increased vertical motion.

From this example and from other investigations on the continuously sampled images, we see that most parameters that are derived directly from the bubble segmentation suffer from a very high sensitivity to froth image perturbations because of the high accuracy of the segmentation that we use. In other words, froth surface images that represent froths that might seem very similar to a trained froth operator yield significantly different extracted parameters.

This is a positive observation, as it is a more correct approach to extract information from accurate parameters that convey physical significance than it is to extract information from under-sensitive parameters which might at first glance seem more directly related to the high-level froth condition and more stable. More simply put, the latter approach could already discard important information at the data-gathering stage, rather than doing so selectively at the decision-making stage. This is very important, especially in the light of the fact that other froth analysis approaches such as the one employed by JKFrothCam [16] deduce froth type from only textural information. It is also in line with the proposed classification approach in [7] where it is stated that images should not be analysed as patterns only, but rather that extracted measures should convey real physical significance.

By employing a mapping function such as a filter or a higher-level classifier, the more accurate parameters can be translated during the decision-making stage to more stable signals representing higher-level characteristics, such as the froth type or condition. It is of primary importance that

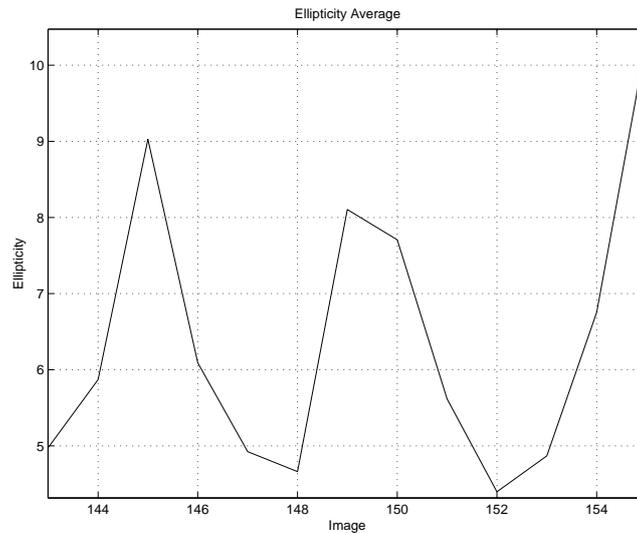


Figure 5.4: Section of ellipticity figure plot for 25 image per second sample set of laboratory froth. See figure 5.2.

the features used for these transforms are as accurate as possible.

This example also serves to demonstrate the detail in which froth bubble dynamics can be examined by using continuous sampling methods. By sampling images at continuous video rates and extracting visual parameters with these algorithms, the physical behaviour of flotation froth can be examined at more depth than possible before.

5.2.2 A Useful Filter for Froth Parameters

Figure 5.5 shows a plot of extracted average bubble area over 400 images of on-plant tank cell froth. The images were captured at 25 images per second and represent 16 seconds of full motion video. Note that even over this relatively short period we see the average bubble area parameter ranging from 490 pixels to a peak of 1950 pixels. The 10-90% interval is at 660 to 1200 pixels, which is still very significant for this short period. The images representing these peaks will shortly be examined.

Note that the extracted parameter plot seems very noisy, although there is a clearly discernible trend. This noise is a complaint often cited by process engineers when faced with the problem of using these signals for control purposes. From investigations of sequences of segmented data, we know that this “noise” is an accurate reflection of occurrences in the captured images. In this case, the “noise” is exacerbated by the fact that the framing of the images includes the flotation cell weir in order to perform motion analysis on froth moving over the cell lip.¹ By the nature

¹Build-up on cell lips can affect recovery. Monitoring mobility over the cell lip is an effective way of monitoring

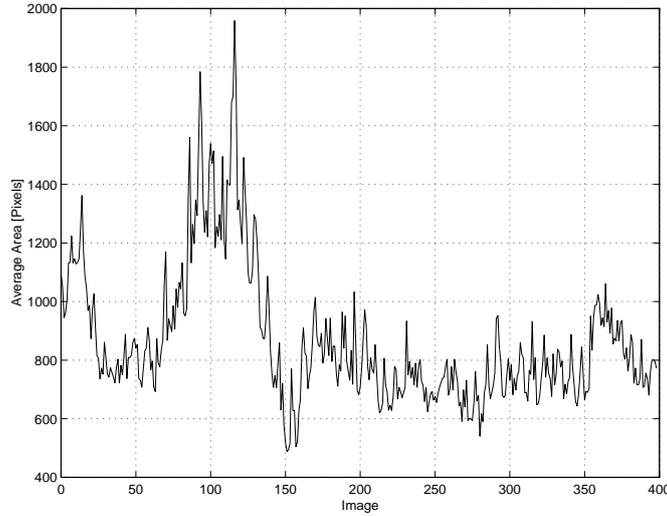


Figure 5.5: Plot of average bubble area/size over 400 images representing 16 seconds of on-plant tank cell froth digitised at 25 images per second.

of its function, bubbles drop away from the camera over the weir, thus causing significantly different segmentations at every image capture and significantly different resultant parameters. The lighting conditions at this point can also be worse as will shortly be seen in the sample images.

A useful filter for extracting the parameter trend is the α -trimmed mean filter [28], which is a member of the class of filters known as L-filters. L-filters are order-statistic and by implication non-linear filters that make a compromise between pure non-linear filtering and pure linear filtering. The output of the L-filter is given by:

$$L(X_1, X_2, \dots, X_N; \mathbf{a}) = \sum_{i=1}^N a_i X_i$$

where X_1, X_2, \dots, X_N are the signal values in a moving window sorted in ascending order of value of size N and

$$\mathbf{a} = (a_1, a_2, \dots, a_N)$$

is a weight vector.

The α -trimmed mean filter is defined as:

$$\text{TrMean}(X_1, X_2, \dots, X_N; \alpha) = \frac{1}{N - 2\alpha N} \sum_{i=\alpha N+1}^{N-\alpha N} X_i$$

where α is the fraction of samples that get trimmed off the start and the end of the sorted moving window. In other words, αN samples are discarded from the beginning and the end of the sorted build-up.

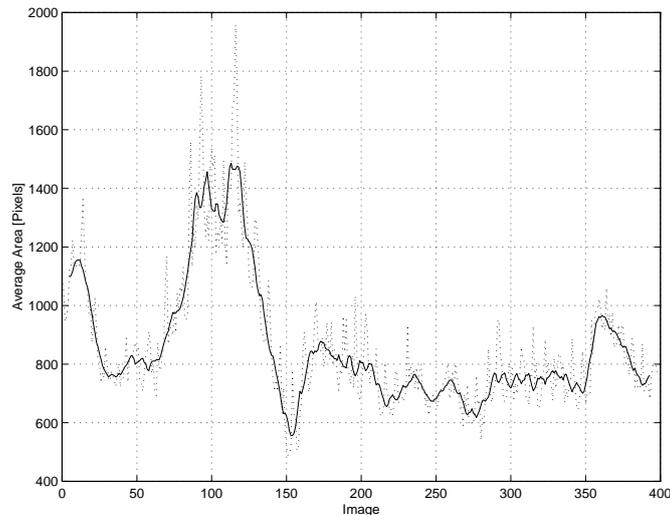


Figure 5.6: Average bubble size signal filtered with α -trimmed mean filter with $N = 10$ and $\alpha = 0.2$

moving window, and the output of the L-filter is the numerical average of the remaining values. Larger α values result in a closer resemblance to a median filter and smaller values result in a closer resemblance to a mean filter.

A median filter discards impulse noise well, but is not well-adapted for additive Gaussian noise removal, whilst the mean filter handles additive Gaussian noise well, but fails badly with impulsive noise [28]. The trimmed mean filter is a compromise which yields good results in cases where impulsive noise and Gaussian noise are present. In our case, there is impulsive “noise” which can be attributed to the sensitivity of the segmentation-based parameters and we assume that the segmentation error can be modelled as Gaussian noise.

Figure 5.6 shows the signal in figure 5.5 after having been filtered with an α -trimmed filter with $N = 10$ and $\alpha = 0.2$.

As a quick illustration of the segmentation performance, the segmentations of images 117, 132 and 152 are shown in figure 5.7. The first image represents the highest peak of the bubble size plot in figure 5.6, the third image represents the deep trough immediately following it, and the second image is halfway between them. From the input images the large drop in average bubble size is clearly visible.

5.2.3 Continuous Video Dynamic Analysis

The average horizontal or vertical motion is the numerical average of the applicable rectangular components of the calculated vector field. The vector field is the array of motion vectors gen-

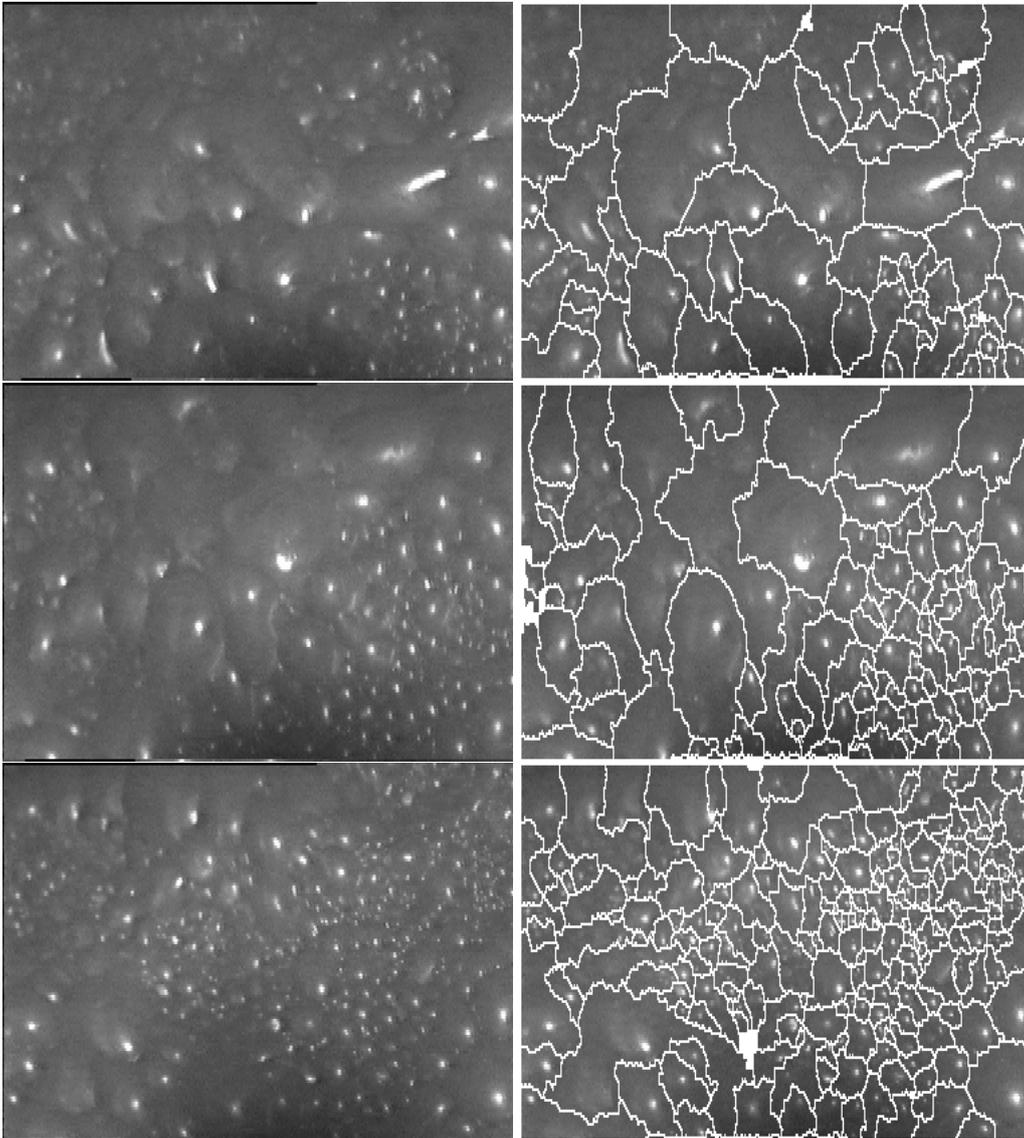


Figure 5.7: Input and segmented froth images corresponding to graph in figure 5.6. Images 117, 132 and 152 from top to bottom.

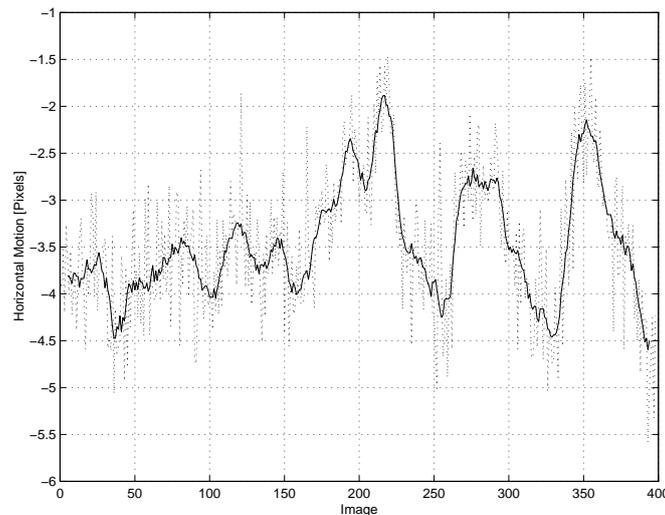


Figure 5.8: Plot of average horizontal pixel motion as calculated by the region-based motion estimation over 400 images sampled at 25 images per second from on-plant scavenger cell. The darker line is a version smoothed with an α -trimmed means filter with $\alpha = 0.2$ and $N = 10$.

erated by the region-based motion estimation, where one motion vector is calculated for each partitioned motion block.

Figure 5.8 is a plot of the region-based motion estimation calculated average horizontal motion over 400 images sampled at 25 images per second from an on-plant scavenger cell and 5.9 is a plot of the per image standard deviation of this parameter. This sampling of the scavenger shows a relatively continuous horizontal flow over 16 seconds and the standard deviation values are relatively small in comparison to the average horizontal motion values that they correspond with. Examining the sample images also show that the motion vectors calculated for each image are generally in full agreement, i.e. they are all of almost equal magnitude and pointing in the same direction. Keep in mind that each vector represents the motion of a separate independent partitioned block of the motion region of interest. This behaviour is to be expected of a flotation cell which is kept under effective control. There are no major step discontinuities in the motion and the standard deviations are relatively small indicating that the region of interest is experiencing constant translational motion.

In order to demonstrate the behaviour of parameters extracted from froth which is not so well-behaved, a sample set of images captured of laboratory laboratory froth was analysed. Figures 5.10 shows a section of the average vertical motion plot and figure 5.11 shows the corresponding per-image standard deviation.

This section has been chosen to illustrate how average motion standard deviation and average motion itself reflect froth events such as bubble bursts, and indirectly froth stability. At image 357 the standard deviation is relatively low in comparison the the vertical motion. The first

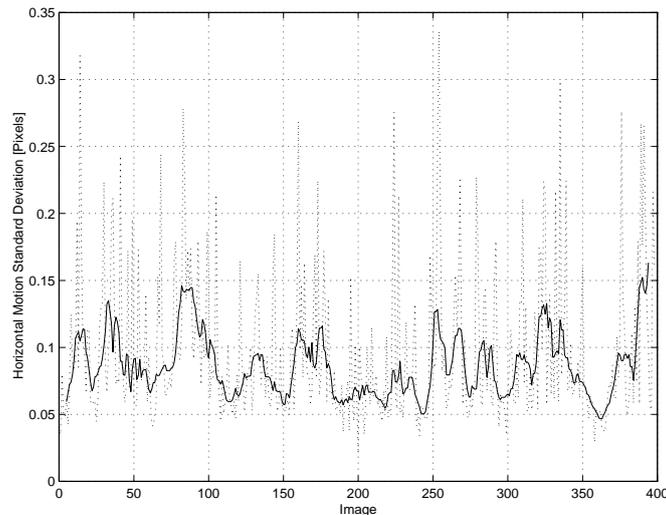


Figure 5.9: Plot of average horizontal pixel motion per image standard deviation as calculated by the region-based motion estimation over 400 images sampled at 25 images per second from on-plant scavenger cell. The darker line is a version smoothed with an α -trimmed means filter with $\alpha = 0.2$ and $N = 10$.

image in figure 5.12 shows that the motion vectors are mostly in agreement about the direction and magnitude of motion. Note that for this experiment, the area-based motion estimation region of interest was set to cover the whole image.

The second image in 5.12, taken 80ms after the first, shows a large bubble that has burst at the bottom center of the image. The motion vectors are in disarray, and this is reflected in the relatively high standard deviation at image 360 (in comparison to the almost zero motion). By image 363 the motion vectors have stabilised and are in agreement again and this is reflected in the standard deviation and vertical motion.

As an aside, the area-based motion estimation is able to detect rotational and deformatory motion when the motion blocks and search regions are relatively small, such as in this example.

5.2.4 Conclusions

Sensitivity of the segmentation-derived parameters is high due to the accuracy of the underlying algorithm, but this is preferable over a stable low-accuracy measurement. The suggested filtering technique for smoothing the segmentation-derived parameters performs well. In practical segmentation-based froth analysis setups at the moment, isolated samples are taken temporally far apart, and static parameters are calculated. Instead of this, short sequences of samples can be taken at each interval and the proposed filtering technique can be used to generate more stable parameters. Short sequences imply between 10 or 20 images at full video sampling rate (i.e. 25

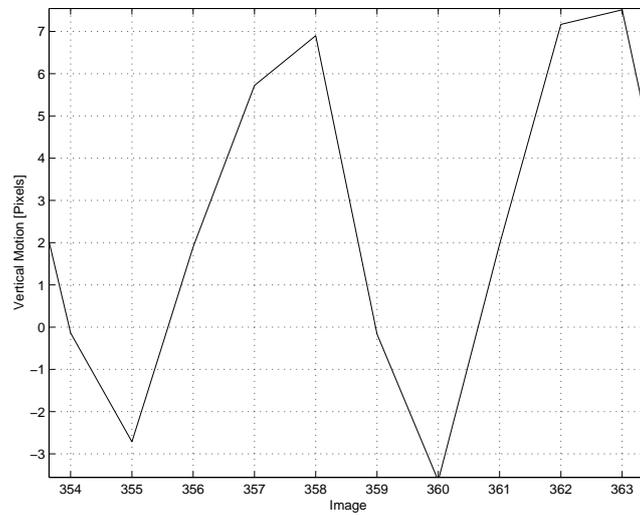


Figure 5.10: Section of plot of average vertical pixel motion for 400 images sampled at 25 images per second from laboratory flotation cell.

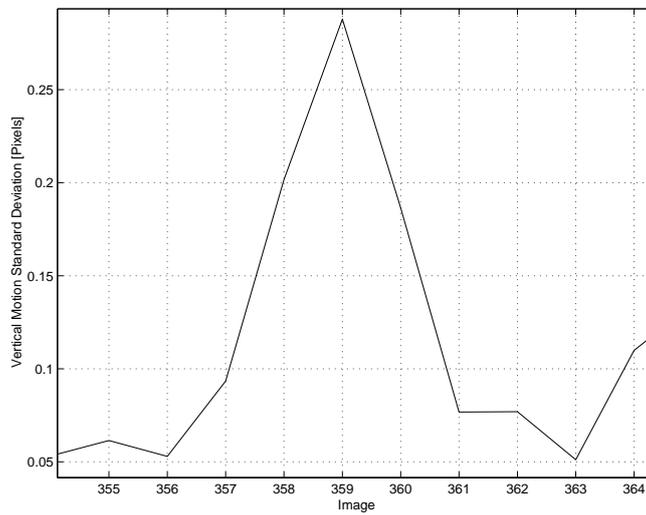


Figure 5.11: Section of plot of average vertical pixel motion per image standard deviation for 400 images sampled at 25 images per second from laboratory flotation cell. Corresponds with figure 5.10

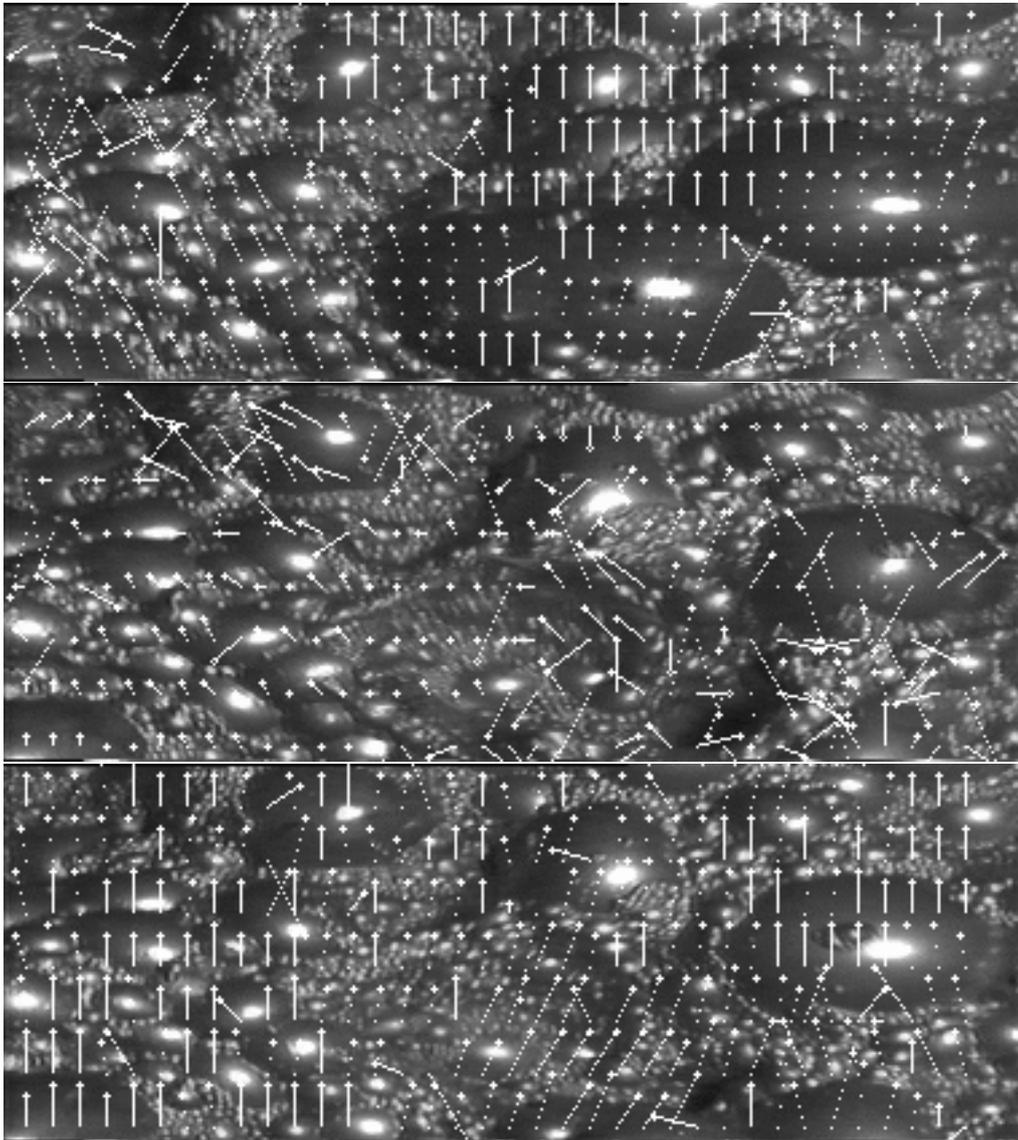


Figure 5.12: Images with motion vectors corresponding with positions 357, 360 and 363 in graphs in figures 5.10 and 5.11.

frames per second), obviously depending on the particular dynamics of the froth under investigation. This would enable the estimation of very accurate yet stable parameters. If more than just stable parameters are desired, i.e. alternative methods for e.g. froth stability quantification, the recommendations in the following paragraphs should be followed.

Continuous video froth analysis can yield more information than the usual image sampling about the flotation process which is being monitored, as temporal relationships between information gathered from successive frames can be utilised for more than just stabilising existing parameters. As an example of this, changes in bubble count between successive images can be used as a more accurate and intuitively pleasing measure of froth stability. As we've seen in this investigation, bubble burst events can cause very large perturbations in the continuous video bubble count parameter. As demonstrated, a mobility analysis where motion and motion standard deviation are taken into account can also be used for stability calculation.

This investigation has shown that although process disturbances have relatively long time constants, there are important froth surface appearance dynamics with lifetimes of as little as tens of milli-seconds. The sampling scheme should be adjusted if these dynamics are to be accounted for.

The summarising recommendation is that the sampling scheme for practical analysis setups should be changed if the full information content of the extracted parameters is to be utilised. Again depending on the particular froth dynamics under investigation, it is recommended that the sampling period is increased drastically, but that 2 or 3 seconds of continuous video or half-rate continuous video (12.5 images per second) is sampled at each sampling interval. In this way, the sampling scheme benefits from the long time constants of the process disturbances themselves, but also detects the low-level froth dynamics, enabling the extraction of accurate and stable parameters as well as the extraction of hitherto unavailable froth information, such as true stability.

5.3 Performance and Impact on Flotation Plants

Recall that froth surface analysis is traditionally performed by human operators who don't report on the number of bubbles they see or the average size, but rather on the deductions they make from their visual observation. These froth operators are judged on the flotation process performance which is a result of the decisions and deductions that they make.

The Fluxar Control System has been running in an advisory fashion on a South African platinum extraction plant since the end of 1998. Only relatively recently have tests started with full closed-loop control by Fluxar, so conclusive results with respects to increased recovery or grade are not yet available. The system is being tested in close-loop control mode every alternative day. On the non-control days it functions in an advisory mode where instead of performing adjustments

automatically, it simply recommends control actions.

Very importantly, reports by process engineers involved at the plant state that during closed-loop control by the Fluxar system the deviation of monitored variables from their ideal set points is smaller than when control is performed by human operators. Also, short-lived deviations which are often missed by human operators are flawlessly handled by the control system. The control system is definitely able to stabilise the flotation process effectively. During advisory operation, the control system's recommendations independently mimic the control actions taken by the froth operators. These things are at the least an added indication that the sensor platform, forming an integral part of the Fluxar control system, is fulfilling its function satisfactorily.

The JKMRC indicate that they have achieved 17% increased recovery of combustible coal at Peak Downs (a coal mine in Australia) and 1.2% at 92% confidence at Escondida (a copper mine in Chile) using their JKFrothCam system [16]. Fluxar-SP is able to extract considerably more visual information than JKFrothCam at high accuracy. It has the potential for the same or better performance than JKFrothCam.

Chapter 6

Conclusions

This chapter presents a synopsis of the research and development done in the course of work. It summarises some of the important discoveries that were made and this work's contribution to the flotation froth analysis industry. It concludes by discussing interesting avenues for future research.

6.1 Summary of Work

The thesis started by giving a short background of the froth flotation process and the necessity and value of employing control, and more specifically automatic control, to maintain process efficiency. It highlighted the advantages of using computational vision as the measurement component of a froth flotation process control system, and then went on to document the relationship of the froth surface appearance to process variables from an image processing perspective. An overview and history of other work on specifically the image processing-based control of the flotation process was given.

Algorithms for the low-level image analysis of flotation froth were investigated. In addition, new methods for detecting fine froth, for increasing the accuracy of the bubble-segmentation and for detecting per-bubble motion were developed. Methods of converting this low-level analysis to parameters describing the physical properties of the froth surface were then discussed, and a novel application of forget-factor smoothing to the calculation of histograms that represent examined parameters over time in a compact way was shown.

As part of this work, a portable and modular implementation of all the froth surface image processing and parameter extraction and post-processing concepts was developed. This froth analysis kernel is an extensive and flexible tool that can be re-used for commercial and research

applications. It is able to perform a full static and two-image dynamic analysis in a few seconds, depending on the capturing and analysis parameters. Literature does not mention a single complete implementation such as this.

Three software packages incorporating this froth analysis kernel were developed. The most extensive one, Fluxar-SP, affords its user easy and user-friendly access to all the froth analysis algorithms and post-processing functionality through a graphical user interface. It is able to acquire image data from four multiplexed video input channels so that it can be deployed in an on-line fashion on flotation plants as the measurement component of an automatic control system. It is also able to acquire data from images stored on disc so that froth images can be analysed off-line. This enables it to be used for after-the-fact analysis of process events as well as a low-cost research tool. Fluxar-SP has proved to be a robust application and has shown in practice that it can fulfil its purpose as the sensor component of an automatic flotation control system. It is in use in selected research setups and in an industrial flotation plant.

Fluxar-Speed is an embedded application of the froth analysis system and is used on plants where a low-cost mobility-monitoring device is required. Because it only has to extract and output mobility information, it can do this at the rate of 10 or more samples per second. This tool is in the process of going to international market.

The Continuous Sampling Platform is a research tool that was designed to capture large batches of continuous video data and analyse them off-line. It performs all available static and dynamic analyses on the captured data and stores all extracted information for experimental post-processing and investigation.

The performance and impact of the algorithms were then investigated. The difficulty of theoretically quantifying froth analysis algorithm accuracy was explained and an attempt was made at partly validating the mobility analyses. A significant preliminary study was done on selected issues pertaining to continuous video flotation froth analysis and it was shown that this is a very promising avenue of research for froth visual analysis. Using the video data that was captured, an important issue of parameter stability vs. parameter accuracy was also discussed. The chapter concluded with a short discussion on the performance and impact of the froth analysis system on flotation plants.

6.2 Contribution

This work yields several important contributions to the froth analysis industry:

1. This work provides a knowledge base for further research and development in computational vision flotation froth analysis specifically, extending the more general knowledge base for machine vision control of flotation provided by Eksteen's work [1].

2. It includes the most complete froth surface visual analysis implementation known to exist. This is in the form of a reusable and portable kernel that can be re-deployed in any number of froth surface analysis applications. In addition, a full static (including complete bubble-segmentation) and dynamic analysis can be performed in a few seconds.
3. Three practical applications have been created. Fluxar-SP is a complete and operator-friendly front-end to the froth analysis kernel for on-plant and research use. Fluxar-Speed is an embedded mobility monitoring device that will find widespread industrial use. The continuous sampling platform is a valuable research tool for investigations into an interesting new direction.
4. Fluxar-SP is installed in a factory and although no conclusive results concerning its performance are available, it is performing very satisfactorily as a robust and accurate froth analysis platform. It is important to note that this development took place in close cooperation with the froth operators and chemical engineers who make use of it, and is thus well-adapted for practical use in industrial environments.
5. Novel additions to froth analysis algorithms:
 - Two methods for fine froth detection in froth images.
 - A heuristic method for post-segmentation over-segmentation correction.
 - An algorithm for the extraction of per-bubble motion.
 - The first application of forget-factor based filtering to froth parameter histograms
6. A preliminary investigation into continuous video froth analysis was performed and selected issues were discussed. There is no literature on previous experiments of this kind. The implementation infrastructure for our investigation can be re-used for further research.

6.3 Future Work

This work opens up several new and interesting opportunities for other research. From an image-processing perspective, a very valuable contribution to the froth surface analysis field would be the creation of a reference database of images that also contains human froth-operator and chemical engineering information describing each image's significance to the process. This database could be used to benchmark froth analysis and classification algorithms, as well as being a very valuable tool in operator training.

The analysis software developed as part of this work facilitates the investigation of various classifiers and other pattern-recognition constructs for the automatic extraction of process variables from imaging parameters. Up to now, most work has been done on classifying froth types on the grounds of texture-related information. Now we have access to parameters representing physical

characteristics of the froth. Performing pattern recognition on these higher-quality features, such as bubble characteristic histograms, promises to be very interesting.

The results of the continuous sampling experiments performed in this work indicate that this is a very promising avenue of research that will help to increase the understanding of automatic froth surface appearance interpretation. Remember that when froth operators utilise their human vision systems to analyse the froth surface and deduce process variables, they are not analysing just one snapshot of the froth surface, but the static and dynamic characteristics over a period of time. Existing automatic froth analysis approaches analyse one or two images at a time.

One of the reasons for this is that the more complex froth analysis methods cannot be performed in real-time at full motion video rates on current processing hardware. This limitation introduces our next opportunity for future research. With the announced commercial availability of mass-produced and affordable 1GHz and over CPUs at the end of this year, the possibility of having complete bubble-segmentation and other complex froth routines execute in real-time on motion video rate froth surface data becomes a conceivable goal. In addition to the improved hardware, this will require further research into additional algorithm optimisation.

One possibility for optimisation could result in a new class of the fundamental algorithms used for froth analysis that rely on the fact that full motion video is being examined and use results on previous images to decrease execution time on subsequent captures. These modifications would have global impact on general computational vision research.

Bibliography

- [1] J. J. Eksteen, “The Development and Technology Transfer of an Industrial Machine Vision System for the control of a Platinum Flotation Plant,” Master’s thesis, University of Stellenbosch, February 1995.
- [2] V. A. Glembotskii, V. I. Klassen, and I. N. Plaksin, *Flotation*. Primary Sources, New York, 1972.
- [3] A. J. Lynch, N. W. Johnson, E. Manlapig, and C. G. Thorne, *Mineral and Coal Flotation Circuits*, vol. 3 of *Developments in Mineral Processing*. Elsevier Scientific Publishing Company, 1981.
- [4] D. W. Moolman, C. Aldrich, G. P. J. Schmitz, and J. S. J. van Deventer, “The interrelationship between surface froth characteristics and industrial flotation performance,” *Minerals Engineering*, vol. 9, no. 8, pp. 837–854, 1996.
- [5] J. Villeneuve, J. C. Guillaneau, and M. V. Durnace, “Flotation modelling: A wide range of solutions for solving industrial problems,” *Minerals Engineering*, vol. 8, no. 4/5, pp. 409–420, 1995.
- [6] M. Bezuidenhout, J. S. J. van Deventer, and D. W. Moolman, “The identification of perturbations in a base metal flotation plant using computer vision of the froth surface,” *Minerals Engineering*, vol. 10, no. 10, pp. 1057–1073, 1997.
- [7] D. W. Moolman, J. J. Eksteen, C. Aldrich, and J. S. J. van Deventer, “The significance of flotation froth appearance for machine vision control,” *International Journal of Mineral Processing*, vol. 48, pp. 135–158, December 1996.
- [8] D. W. Moolman, C. Aldrich, J. S. J. van Deventer, and W. W. Stange, “The classification of froth structures in a copper flotation plant by means of a neural net,” *International Journal of Mineral Processing*, vol. 43, pp. 193–208, 1995.
- [9] D. W. Moolman, C. Aldrich, J. S. J. van Deventer, and D. J. Bradshaw, “The interpretation of flotation froth surfaces by using digital image analysis and neural networks,” *Chemical Engineering Science*, vol. 50, no. 22, pp. 3501–3513, 1995.

- [10] P. J. Symonds and G. de Jager, "A technique for automatically segmenting images of the surface froth structures that are prevalent in industrial flotation cells," in *Proceedings of the South African Symposium on Communications and Signal Processing*, pp. 111–115, 1992.
- [11] D. W. Moolman, C. Aldrich, and J. S. J. van Deventer, "The monitoring of froth surfaces on industrial flotation plants using connectionist image processing techniques," *Minerals Engineering*, vol. 8, no. 1–2, pp. 23–30, 1995.
- [12] L. H. Siew, R. M. Hodgson, and E. J. Wood, "Texture Measures for Carpet Wear Assessment," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, pp. 92–105, January 1988.
- [13] N. Sadr-Kazemi and J. J. Cilliers, "An image processing algorithm for measurement of flotation froth bubble size and shape distributions," *Minerals Engineering*, vol. 10, no. 10, pp. 1075–1083, 1997.
- [14] J. J. Francis and G. de Jager, "An investigation into the suitability of various motion estimation algorithms for froth imaging," in *Proceedings of the IEEE South African Symposium on Communications and Signal Processing*, pp. 139–142, September 1998.
- [15] A. Cipriano, M. Guarini, R. Vidal, A. Soto, C. Sepúlveda, D. Mery, and H. Briseño, "A real time visual sensor for supervision of flotation cells," *Minerals Engineering*, July–August 1997.
- [16] P. N. Holtham, K. K. Nguyen, and P. M. Cameron, "Recent Developments in the Application of Froth Imaging Techniques to the Control of Flotation Plants," in *Mineral Processing*, South African Institute of Mining and Metallurgy, 1999.
- [17] L. Vincent and P. Soille, "Watersheds in digital spaces: An efficient algorithm based on immersion simulations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, pp. 583–598, June 1991.
- [18] L. Vincent, "Morphological grayscale reconstruction in image analysis: Applications and efficient algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 2, pp. 176–201, April 1993.
- [19] E. R. Dougherty, *An Introduction to Morphological Image Processing*. SPIE Optical Engineering Press, 1992.
- [20] C. P. Botha, D. M. Weber, M. van Olst, and D. W. Moolman, "A practical system for real-time on-plant flotation froth visual parameter extraction," in *Proceedings of IEEE Africon*, September 1999.
- [21] S. Zafar, Y.-Q. Zhang, and J. S. Baras, "Predictive Block-Matching Motion Estimation for TV Coding – Part I: Inter-Block Prediction," *IEEE Transactions on Broadcasting*, vol. 37, pp. 97–101, September 1991.

-
- [22] A. K. Jain, *Fundamentals of Digital Image Processing*. Prentice Hall, 1989.
- [23] R. N. Bracewell, *Two-Dimensional Imaging*. Prentice Hall, 1995.
- [24] C. Stiller and J. Konrad, "Estimating Motion in Image Sequences," *IEEE Signal Processing Magazine*, vol. 16, pp. 70–91, July 1999.
- [25] J. Peyton Z. Peebles, *Probability, Random Variables, and Random Signal Principles*. McGraw-Hill, 1993.
- [26] M. Nadler and E. P. Smith, *Pattern Recognition Engineering*. John Wiley & Sons, 1992.
- [27] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis and Machine Vision*. Chapman & Hall, 1993.
- [28] J. Astola and P. Kuosmanen, *Fundamentals of Nonlinear Digital Filtering*. CRC Press, 1997.
- [29] Crusader Systems, "Crusader Systems Company Web Site," 1999. <http://www.crusader.co.za/>.
- [30] Crusader Systems, "Fluxar Web Site," 1999. <http://www.fluxar.com/>.
- [31] K. V. Mardia, *Statistics of Directional Data*. Academic Press, 1972.

Appendix A

Circular Statistics

This appendix summarises the findings with respect to basic statistic calculations on circular data. For further information on spherical (and higher dimensional directional) data, consult the references [31].

A.1 Selected statistic measures

A.1.1 Mean Direction

The mean direction is defined to be the direction of the resultant of the unit vectors representing the sampled angles¹, in other words the mean direction \bar{x}_0 is defined as

$$\bar{x}_0 = \arctan\left(\frac{\sum_{i=1}^n \sin \theta_i}{\sum_{i=1}^n \cos \theta_i}\right)$$

with θ_i being the i 'th angle sample and n being the number of samples.

This definition of the mean direction is independent of the choice of zero-axis, which is a desirable property. Also, the sum of deviations about the mean equals zero, i.e.

$$\sum_{i=1}^n \sin(\theta_i - \bar{x}_0) = 0$$

¹I am obliged to point out that Dave was right so as to avoid being assaulted by repeated “I told you so”-type remarks from the aforementioned.

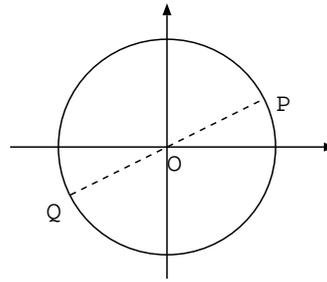


Figure A.1: Illustration of median direction

A.1.2 Median Direction

The median direction ζ_0 of a set of sample angles is any solution of

$$\int_{\zeta_0}^{\zeta_0+\pi} f(\theta) d\theta = \int_{\zeta_0+\pi}^{\zeta_0+2\pi} f(\theta) d\theta = \frac{1}{2}$$

where $f(\theta)$ satisfies $f(\zeta_0) > f(\zeta_0 + \pi)$.

In more graphic terms, if all the angles in the sample set are considered as points P_i on a unit circle, any one of these, P , is considered a median if the line PQ (with Q being diametrically opposite P) has an equal number of sample points on either side, and the majority of the sample points are closer to P . See figure A.1.

For data on a line, it is sufficient to perform sorting and pick the middle element. In the circular case, it is easy to see that this can yield incorrect results. For example, consider an example from our references with sample angles 43, 45, 52, 61, 75, 88, 88, 279 and 357 (in degrees) where the median direction according to the above definition is 52, whilst for the case on the line it would've been 75.

Note that for data falling in the interval 0 to 180° , the above definition simplifies to sorting and selecting the middle element.

A.1.3 Circular Variance and Standard Deviation

Again representing sample angles with points on the unit circle, consider P representing the zero direction, and P_i a sample point, then a measure of dispersion ζ_i of the sample point from the point representing the zero direction could be defined as

$$\zeta_i = \min(\theta_i, 2\pi - \theta_i) = \pi - |\pi - \theta_i|$$

i.e. the smallest angle between the two vectors \overline{OP} and $\overline{OP_i}$.

A measure of total dispersion D of all sample points from P could then be defined as

$$D = \frac{1}{n} \sum_{i=1}^n (1 - \cos \zeta_i)$$

and

$$D = \frac{1}{n} \sum_{i=1}^n (1 - \cos(\theta_i - \alpha)) \quad (\text{A.1})$$

if we rotate the zero direction by α degrees, i.e. we are measuring the total dispersion from the vector representing an angle of α degrees.

Setting equation A.1 equal to zero and taking the derivative with respect to α yields

$$\sum_{i=1}^n \sin(\theta_i - \alpha) = 0$$

showing that the dispersion is smallest about α . If $\alpha = \bar{x}_0$ (i.e. the mean) then we have that the dispersion (equation A.1) is a minimum about the mean, which is a desirable property.

We can write D as a dispersion about the mean \bar{x}_0 and call it S_0 , the circular variance, as in

$$S_0 = 1 - \frac{1}{n} \sum_{i=1}^n \cos(\theta_i - \bar{x}_0) \quad (\text{A.2})$$

Setting

$$\bar{R} = \frac{1}{n} \sqrt{\left(\sum_{i=1}^n \cos \theta_i\right)^2 + \left(\sum_{i=1}^n \sin \theta_i\right)^2}$$

with \bar{R} being the mean resultant length (of all the unit vectors representing the sample angles) and keeping in mind that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \cos \theta_i &= \bar{R} \cos \bar{x}_0 \\ \frac{1}{n} \sum_{i=1}^n \sin \theta_i &= \bar{R} \sin \bar{x}_0 \end{aligned}$$

equation A.2 can be reduced to

$$S_0 = 1 - \bar{R}$$

In order to transform this to the interval $(0, \infty)$, the following transformation is applicable:

$$s_0 = \sqrt{\{-2 \ln(1 - S_0)\}}$$

which is defined as the circular standard deviation.

A.2 Angular Distributions on the range $(0, 2\pi/l)$

Consider a random variable θ that is entirely on the range $(0, 2\pi/l)$, i.e. regard it as only having values in that range. Define

$$\theta' = l\theta$$

which will have values on the interval $(0, 2\pi)$.

Now, if we define \bar{x}'_0 and s'_0 respectively as the mean direction and circular standard deviation of θ' , then the mean of θ

$$\bar{x}_0 = \bar{x}'_0/l$$

and the deviation of θ

$$s_0 = s'_0/l$$

Defining the translated variance is not as easy. For a distribution concentrated in a small arc about the mean, the relation

$$V_0 = V'_0/l^2$$

works well. In other cases, for wrapped normal and Von Mises distributions (see the reference)

$$V_o = 1 - (1 - V'_0)^{1/l^2}$$

is a valid general relation.

For the median direction, the method documented in section [A.1.2](#) can be used as is, no angle translation has to be done.