

There is no quality in this world that is not what it is merely by contrast.
Nothing exists in itself.

Herman Melville (1819-1891)

University of Alberta

CONTRAST ENHANCEMENT OF OIL SAND IMAGES USING MORPHOLOGICAL SCALE SPACE

by

Andrzej Zadorożny



A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of **Master of Science**.

Department of Computing Science

Edmonton, Alberta
Spring 2006



Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file *Votre référence*

ISBN: 0-494-13918-8

Our file *Notre référence*

ISBN: 0-494-13918-8

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

To my Mom and Dad.
Without you nothing would be possible.

Abstract

In this thesis, we present a new algorithm of contrast enhancement of oil sand images. Image contrast enhancement constitutes a crucial pre-processing step in the segmentation of such images, used for granulometry purposes. Oil sand images are difficult to segment due to their complex nature. The new contrast enhancement method aims at improving the quality of the images in order to improve segmentation performance. The method uses a multi-scale image decomposition, obtained with a series of morphological top-hat transformations, where the scale of enhancement corresponds to expected object size. In addition, the new enhancement method is direct, where the level of enhancement is controlled based on a contrast measure. Also, the new method is adaptive, where the enhancement is applied locally, based on local image properties. We present experimental results, where we quantitatively measure the improvement in the quality of oil sand image segmentation, after applying our contrast enhancement method.

Acknowledgements

I would like to express my sincere gratitude to those who have supported this research and helped me along the way. Without their help the work presented here would not be possible.

I owe a great deal of thanks to my supervisor, Dr. Hong Zhang, who has been an incredible mentor and a friend through all these years. He has greatly encouraged and supported me during all this time, always opening new doors, and always promoting new frontiers of exploration.

I am very grateful to the members of my committee, Dr. Hong Zhang, Dr. Xiaobo Li, Dr. Ron Kube and Dr. Benoit Rivard, who took the time to read and comment on this work.

Special thanks go to Dr. Ron Kube for his great advice and support of the industrial end of this research, and for a lot of valuable research and writing tips.

I would like to thank NSERC and Syncrude Canada Ltd. for their financial support of this research, and I would also like to acknowledge Syncrude Canada Ltd. for providing the images of oil sands.

Big thanks go to Dr. Mark Polak and Dr. Minghong Pi for their collaboration on the scoring metric, and also to Dr. Mark Polak for providing the fragment size distribution code. Also big thanks to David Laing for his work on OSA and the tuning of OSA's parameters. I would like to thank Dr. Xiujian Luo for generating the ground-truth segmentation images. Also, thanks to Dr. Ricardo J. Ferrari for the ideas on energy distributions in multi-scale representations. My gratitude also goes to Luca Pireddu for his tips on thesis writing and to Sean Verret for providing the LaTeX templates for this thesis.

Finally I would like to thank all those whom I missed here, but who also helped and supported me throughout this research.

Table of Contents

1	Introduction	1
1.1	The Problem	1
1.2	Contrast Enhancement Approaches	3
1.3	Thesis Objectives and Contributions	4
1.4	Methodology and Results	4
1.5	Overview	4
2	Background and Related Work	6
2.1	Introduction	6
2.2	Digital Image Processing	6
2.3	Digital Image Enhancement	7
2.4	Contrast Enhancement	8
2.5	Classification of Contrast Enhancement Methods	8
2.5.1	Global vs. Local Methods	8
2.5.2	Direct vs. Indirect Methods	9
2.6	Different Approaches of Enhancing Contrast	10
2.6.1	Contrast Stretching	10
2.6.2	Full Frame Histogram Modification	11
2.6.3	Local Area Histogram Modification	12
2.6.4	Unsharp Masking	15
2.6.5	Contrast Enhancement in Frequency Domain	17
2.6.6	Multi-Scale Enhancement	17
2.7	Summary	20
3	Image-Based Granulometry of Hard Rocks and Oil Sands	22
3.1	Introduction	22
3.2	Image-Based Granulometry	22
3.3	Granulometry of Hard Rocks	24
3.3.1	Measuring Methods	24
3.3.2	Image-Based Measuring Systems	25
3.3.3	Split [®]	25
3.3.4	WipFrag [®]	26
3.4	Granulometry of Oil Sands	26
3.5	Current Segmentation Method of Oil Sand Images	27
3.5.1	Morphological Segmentation	27
3.5.2	Ore Size Analyst	28
3.5.3	Contrast Enhancement in OSA	28
3.6	Summary	29

4	Morphological Multi-Scale Contrast Enhancement	31
4.1	Introduction	31
4.2	Basic Assumptions	31
4.3	Top-Hat Transformation	32
4.4	Top-Hat Multi-Scale Decomposition	35
4.5	Properties of the Multi-Scale Decomposition	36
4.6	Motivation	40
4.7	A Priori Knowledge	40
4.8	Contrast Measure	42
4.9	Control Value	43
4.10	Image Enhancement	45
4.11	Parameters and Their Tuning	47
4.12	Fast Computation of the Top-Hat Transformation	47
4.13	Summary	49
5	Experiments and Results	52
5.1	Introduction	52
5.2	Ground-Truth Images	52
5.3	Evaluation using Scoring Metric	53
5.4	Experimental Setup	55
5.5	Results	57
5.6	Discussion	63
5.7	Summary	63
6	Conclusions	65
6.1	Summary	65
6.2	Limitations	66
6.3	Future Work	66
	Bibliography	68
A	Individual Experiment Results	70
A.1	Result Images and Plots	70

List of Tables

4.1	Empirically obtained parameters that are used in MMS.	47
5.1	Segmentation conditions handled by the score metric.	55
5.2	Explanations of abbreviations used in Figure 5.2.	56
5.3	Numbers of segmented fragments for the first 6 iterations ($1 \leq i \leq 6$) of MMS. The numbers of fragments of the initial OSA iteration, as well as using the ground-truth as a priori data are also included.	59
5.4	Results of the experiment for the first 6 iterations ($1 \leq i \leq 6$) of MMS. The results of the initial OSA iteration, as well as using the ground-truth as a priori data are also included.	59
5.5	Absolute improvement, derived from Table 5.4.	60
5.6	Relative improvement, derived from Table 5.4.	60
5.7	Potential captured, derived from Table 5.4.	60

List of Figures

1.1	An example image showing oil sand fragments. The individual objects have very complex texture, and their intensity varies significantly. In many cases the edge information is poor and not easily distinguishable.	2
1.2	An example image of hard rock fragments. The fragments in this image are clearly visible with strong edge definition and relatively uniform texture and intensity.	3
2.1	Many direct contrast enhancement methods use a local contrast measure that deals with computing the relationship between a foreground pixel and its local contextual background.	9
2.2	Examples of contrast stretching transformation functions (redrawn from [10]).	11
2.3	An example image of crushed oil sand fragments on a conveyor belt and corresponding enhanced images via global histogram equalization and local histogram equalization. The histograms of the three images are also included. Notice how the histogram of the image enhanced by the global method is stretched, whereas the histogram of the image enhanced by the local method has more uniform distribution.	13
2.4	Contextual region in Gordon's method.	14
2.5	An example of the contextual region in Adaptive Neighbourhood Contrast Enhancement (ANHE).	14
2.6	An example of the basic ANECE method and its power variation version, enhancing the image from Figure 2.3(a). The histograms of the two images are also included. Notice how the power variation version results in a smoother histogram, decreasing the effect of noise over-enhancement.	16
2.7	An example of contrast enhancement using non-linear RoLP pyramid recombination, enhancing the image from Figure 2.3(a). Image histogram is also included. This histogram has a bi-modal characteristic. In addition, bright features are overemphasized.	18
2.8	An example of contrast enhancement using multi-scale top-hat decomposition, enhancing the image from Figure 2.3(a). The image, due to emphasis of fine features, is sharper, with significant enhancement of texture and noise. Image histogram is also included. The peak at grey value 0 indicates considerable enhancement of black features.	20
3.1	An image of oil sand ore on a conveyor belt.	23
3.2	Typical steps in image-based granulometry	24
3.3	The setup for image-based oil sand granulometry.	27
4.1	An illustration of the greyscale white top-hat transformation.	33
4.2	An illustration of the greyscale black top-hat transformation (also known as bottom-hat transformation).	33
4.3	An example of a white and a black top-hat transformation on the image from Figure 3.1.	34

4.4	The top-hat multi-scale decomposition has two independent components, each consisting of decomposed images at different scales. The input image can be exactly reconstructed from any of the two components.	35
4.5	An example of base images at different scales of the white-feature top-hat multi-scale decomposition of the image in Figure 3.1. All images are enhanced for viewing purposes.	36
4.6	An example of scale slice images at different scales of the white-feature top-hat multi-scale decomposition of the image in Figure 3.1. All images are enhanced for viewing purposes.	37
4.7	Artificial samples with disks of different sizes, and corresponding normalized energy distributions. Notice how the peak in the energy response corresponds with the disk radius.	38
4.8	Image from Figure 4.7(c), using disks with radius $r = 12$, with various distortions and corresponding normalized energy distributions. The Gaussian blur is obtained via convolution with a Gaussian kernel with $\sigma = 3$. The noise is generated by adding a random value in the range $[-a, a)$ to each pixel, where $a = 10$	39
4.9	The control value function is based on the Gaussian function.	44
4.10	The control value function depends on the expected scale of objects.	44
4.11	Pseudo code of the MMS algorithm.	45
4.12	An example illustrating the implementation of shifting of a structuring element. During shifting only edge pixels marked with “+” and “-” need to be considered, and all other pixels are ignored. Pixels marked with “+” are added to a local histogram, and pixels marked with “-” are subtracted.	48
4.13	An example of a histogram under the structuring element, illustrating the implementation of morphological erosion and dilation in a single operation. In this operation both the minimum and the maximum values are tracked simultaneously.	49
4.14	An example illustrating the implementation of morphological opening and closing in a single operation. The erosion in the opening and the dilation in the closing both take the same image as input, thus they can be combined into a single erosion/dilation operation.	50
5.1	An example of an oil sand input image and its corresponding manual ground-truth segmentation.	53
5.2	Experimental setup for testing of the MMS algorithm. Abbreviations used in this figure are explained in Table 5.2.	56
5.3	Example results for Image A, from Figure 5.1(a). This examples includes the initial iteration of OSA, and the first and sixth iteration of MMS. Corresponding segmented images are also shown. A complete set of result images for all 10 input images is included in Appendix A.	57
5.4	Example results for Image A, from Figure 5.1(a). This examples includes the initial iteration of OSA, and the first and sixth iteration of MMS. Corresponding image histograms are also shown. The histograms from the MMS method have a bi-modal characteristic, showing the separation between the bright and the dark features. Subsequent iterations of MMS result in a smoother histogram. MMS enhances some pixels to the maximum grey value, which can be seen by the peaks at grey value 255 in the corresponding histograms.	58
5.5	Results for the first 50 iterations of average score for Images A-J. The most improvement is produced by the first iteration of OSA. The maximum average segmentation score of 0.6335 is observed in iteration 6 of MMS. In the subsequent iterations, the average segmentation score is lower, but remains relatively uniform. Individual plots for each of the 10 test images are included in Appendix A.	62

A.1	Enhancement and segmentation results for Image A.	71
A.2	Enhancement and segmentation results for Image B.	72
A.3	Enhancement and segmentation results for Image C.	73
A.4	Enhancement and segmentation results for Image D.	74
A.5	Enhancement and segmentation results for Image E.	75
A.6	Enhancement and segmentation results for Image F.	76
A.7	Enhancement and segmentation results for Image G.	77
A.8	Enhancement and segmentation results for Image H.	78
A.9	Enhancement and segmentation results for Image I.	79
A.10	Enhancement and segmentation results for Image J.	80
A.11	Cumulative percentage passing for Image A (213 fragments in ground truth).	81
A.12	Cumulative percentage passing for Image B (171 fragments in ground truth).	82
A.13	Cumulative percentage passing for Image C (237 fragments in ground truth).	83
A.14	Cumulative percentage passing for Image D (213 fragments in ground truth).	84
A.15	Cumulative percentage passing for Image E (303 fragments in ground truth).	85
A.16	Cumulative percentage passing for Image F (171 fragments in ground truth).	86
A.17	Cumulative percentage passing for Image G (207 fragments in ground truth).	87
A.18	Cumulative percentage passing for Image H (213 fragments in ground truth).	88
A.19	Cumulative percentage passing for Image I (273 fragments in ground truth).	89
A.20	Cumulative percentage passing for Image J (309 fragments in ground truth).	90
A.21	Results of the first 50 iterations of MMS on Image A.	91
A.22	Results of the first 50 iterations of MMS on Image B.	92
A.23	Results of the first 50 iterations of MMS on Image C.	93
A.24	Results of the first 50 iterations of MMS on Image D.	94
A.25	Results of the first 50 iterations of MMS on Image E.	95
A.26	Results of the first 50 iterations of MMS on Image F.	96
A.27	Results of the first 50 iterations of MMS on Image G.	97
A.28	Results of the first 50 iterations of MMS on Image H.	98
A.29	Results of the first 50 iterations of MMS on Image I.	99
A.30	Results of the first 50 iterations of MMS on Image J.	100

List of Abbreviations

Abbreviation	Description
ACE	Adaptive Contrast Enhancement
AHE	Adaptive Histogram Equalization (same as LAHE)
ANECE	Adaptive-Neighbourhood Extended Contrast Enhancement
ANHE	Adaptive-Neighbourhood Histogram Equalization
CIMS	Centre for Intelligent Mining Systems
CLAHE	Contrast-Limited Adaptive Histogram Equalization
DCT	Discrete Cosine Transform
JPEG	Joint Photographic Experts Group
LAHE	Local-Area Histogram Equalization (same as AHE)
MMS	Morphological Multi-Scale (contrast enhancement)
OSA	Ore Size Analyst
RoLP	Ratio of Low Pass

List of Equations

Equation	Description
(2.1)	An example of a contrast measure
(2.2)	Enhancement by Wallis and Lee [19]
(2.3)	Contrast measure by Gordon and Rangayyan [11]
(2.4)	Enhancement of the contrast measure from Equation (2.3)
(2.5)	Enhanced pixel from Equation (2.3), for $p \geq a$
(2.6)	Enhanced pixel from Equation (2.3), for $p < a$
(2.7)	A transformation function [10]
(2.8)	Contrast measure in the ACE algorithm [2]
(2.9)	Mean edge grey level in the ACE algorithm
(2.10)	Edge value in Equation 2.11
(2.11)	Enhancement in the ACE algorithm
(2.12)	Linear unsharp masking [10]
(2.13)	Contrast measure in the Tang's method [31]
(2.14)	Contrast enhancement by the Tang's method
(2.15)	H_n used in Equation (2.14)
(2.16)	White top-hat transformation by Mukhopadhyay and Chanda [20]
(2.17)	Contrast enhancement by Mukhopadhyay and Chanda
(2.18)	White features image
(2.19)	Black features image
(3.1)	Rosin-Rammler material distribution [3]
(4.1)	Greyscale morphological erosion
(4.2)	Greyscale morphological dilation
(4.3)	Greyscale morphological opening
(4.4)	Greyscale morphological closing
(4.5)	White top-hat transformation
(4.6)	Black top-hat transformation (also known as bottom-hat transformation)
(4.7)	Radius of the equivalent circle
(4.8)	Adjusted scale map
(4.9)	Contrast measure in MMS
(4.10)	Inverse transformation of Equation (4.9)
(4.11)	Gaussian function
(4.12)	Control value for enhancing white features
(4.13)	Control value for enhancing black features
(4.14)	Enhancement of the contrast measure for white features
(4.15)	Enhancement of the contrast measure for black features
(4.15)	Enhancement of the contrast measure for black features
(4.16)	Composition of the final enhanced image

Equation	Description
(5.1)	Intersection over union
(5.2)	$\delta(x)$ used by Equation (5.3)
(5.3)	Partial score metric
(5.4)	Final score metric
(5.5)	Definition of an input binary image used
(5.6)	Conditions for score metric value of 1
(5.7)	Conditions for score metric value of 0
(5.8)	Symmetric property of the score metric
(5.9)	Scale invariant property of the score metric
(5.10)	Potential captured

Chapter 1

Introduction

When we look at images, we often think about the quality of those images. What makes an image have a good quality? Why is that quality good? To answer those questions, we often need to focus on the purpose of such images. Are those images intended for appreciation of art, or perhaps for an examination by a professional, such as a radiologist? Maybe they are not intended for a human visual system at all, but instead are an input to a computer algorithm that performs further processing on these images. Whatever the purpose is, the definition of quality of such images goes along with that purpose.

One way to evaluate the quality of images is to look at the tasks that need to be performed on those images, whether intended for humans or for computers. By focusing on the tasks, we can ask ourselves what aspects of those images make the tasks easier. One such aspect is image contrast. It is the contrast that determines how easy it is to perceive the information in the images, and how easy it is to distinguish various details in the image, locating objects of interest.

Some example tasks performed on images fall under the category of image segmentation. Whether performed by human vision or by computer algorithms, image segmentation deals with dividing an image into various regions or objects. Image contrast, being defined as a difference in intensity, has an influence on how easy it is to perform the tasks of distinguishing various objects or features in an image. Image contrast has a significant influence on the quality of the image.

1.1 The Problem

Oil sand is a composition of sand, mineral rich clays, water and bitumen. The oil in oil sands is found in bitumen, which in its raw state is a black, thick, asphalt-like substance. In order to be transportable by pipelines and usable by refineries, bitumen undergoes a process of upgrading. The upgraded bitumen is called sweet crude oil, and it consists of a mixture of light and heavy gas oils referred to as naphtha. Oil sand, in its composition, contains

about four percent of water, which surrounds each grain of sand, and separates bitumen from the sand. This layer of water allows the extraction of oil by water-based methods [17].

The problem we discuss in this thesis is the enhancement of contrast in images of oil sands for the applications of image-based granulometry. Oil sand ore is crushed, breaking the ore into smaller pieces, which we refer to as *fragments*, *objects*, *particles* or *rocks*. Images containing oil sand fragments on a conveyor belt are segmented in order to determine the distribution of the sizes of those fragments. Contrast enhancement is a crucial pre-processing step in the segmentation of such images.




Figure 1.1: An example image showing oil sand fragments. The individual objects have very complex texture, and their intensity varies significantly. In many cases the edge information is poor and not easily distinguishable.

The nature of oil sand images makes accurate segmentation a very difficult task. The appearance of oil sand is affected by the variance in oil sand composition. The objects in oil sand images have a very complex texture. Their intensity often varies from object to object as well as within the individual objects. In many cases bright patches of clay are mixed with dark bitumen. In addition the fragments in oil sand images have poorly defined edges, which is often caused by the softness of the material. Oil sands are usually moist and covered with fine particles. An example of an image of oil sand particles on a conveyor belt is shown in Figure 1.1.

Much literature that exists on measurement of rock fragmentation deals with the hard rock industry. Images of hard rocks, as opposed to oil sands, are much easier to segment. They have clearly defined edges, uniform texture, and the intensity of the individual objects does not vary. An example of such an image is shown in Figure 1.2. Often applications in the fragmentation measurements of hard rock use edge-based techniques in the segmentation of those images. In the case of oil sands, edge-based techniques fail.



Figure 1.2: An example image¹ of hard rock fragments. The fragments in this image are clearly visible with strong edge definition, and relatively uniform texture and intensity.

1.2 Contrast Enhancement Approaches

There are several contrast enhancement methods. Some methods are implemented in spatial domain, and some in frequency domain. Also, contrast enhancement approaches can be local or global, as well as direct or indirect. In addition some methods are based on a single scale processing, while others process multi-scale image representations.

Examples of classical methods include contrast stretching [10], histogram equalization [12] and unsharp masking [28]. Some methods are derived from those classical approaches, such as contrast-limited adaptive histogram equalization (CLAHE) [25], adaptive neighbourhood histogram equalization (ANHE) [22] and adaptive unsharp masking [28].

Direct contrast methods define and evaluate a contrast measure while enhancing the images. Examples of those methods are Gordon's method [11], adaptive contrast enhancement (ACE) [2] and adaptive neighbourhood extended contrast enhancement (ANECE) [19].

Other contrast enhancement approaches include the use of frequency domain and multi-scale decompositions. Examples of these methods include the Tang's method [31], which operates in the discrete cosine transform (DCT) domain and wavelet-based approaches, such as Jin's method [15], which applies LAHE to individual frequency bands. In addition there are also methods based on multi-scale decompositions. These include Toet's method [32] that uses ratio of low-pass pyramid, and a method by Mukhopadhyay and Chanda [20], which uses morphological top-hat transformations. The contrast enhancement methods are discussed in further detail in Chapter 2.

¹The image in Figure 1.2 has been obtained from the Split Engineering LLC website at <http://www.spliteng.com/> and included in this thesis by special permission.

1.3 Thesis Objectives and Contributions

The work presented in this thesis focuses on contrast enhancement of images of oil sands, as an essential pre-processing step in the segmentation of such images. We test the hypothesis that multi-scale contrast enhancement, in which the scale of enhancement corresponds to the size of the object being enhanced, improves the accuracy of oil sand image segmentation.

This thesis makes a number of research contributions. We develop a new contrast-enhancement algorithm for improving the quality of oil sand images. We call this algorithm morphological multi-scale (MMS) contrast enhancement. The new algorithm is adaptive, multi-scale and direct. In addition, we demonstrate that the scale of enhancement is related to the size of the object being enhanced. We also show that better segmentation results are achieved with the improvement of the contrast of input images.

1.4 Methodology and Results

The MMS algorithm that we develop for the enhancement of oil sand images is based on a morphological top-hat multi-scale decomposition. This decomposition is based on a series of top-hat transformations obtained using morphological operations of opening and closing. The scale in the top-hat decomposition is controlled by the size of the structuring element. This algorithm performs local enhancement on each scale of the decomposition. In addition this method is direct, in which the amount of contrast enhancement is controlled via a local contrast measure.

The MMS method takes advantage of a priori knowledge of estimated object sizes. This a priori knowledge provides an initial state, where as a result of the contrast enhancement, subsequent segmentation steps are able to refine this estimation. Expected size of each given object corresponds to the range of scales, on which the object is enhanced.

We evaluate the new contrast-enhancement algorithm quantitatively in terms of segmentation accuracy that results when using this algorithm. The segmented images are compared against ten manually generated ground-truth images, and a segmentation score calculated based on how well the segmented images match the provided ground-truth. The experimental results that we obtain show that the new algorithm improves the accuracy of oil sand image segmentation. The improvement is observed in all ten cases, which makes the results statistically significant.

1.5 Overview

The rest of this thesis is presented in the chapters that follow. Chapter 2 talks about the background and related work in the field of image contrast enhancement. Chapter 3 discusses granulometry of hard rocks and oil sand images, pointing out numerous challenges

with the segmentation of oil sand images. In Chapter 4 we discuss the details of the new contrast enhancement algorithm, and the concepts behind. In Chapter 5 we describe experiments performed in order to evaluate the new algorithm, and present experimental results. Finally in Chapter 6 we provide the conclusions of this thesis, including a discussion on the contributions learned and suggestions for future work.

Chapter 2

Background and Related Work

The previous chapter introduces various concepts to be discussed in this thesis. This chapter presents background and related work in the area of contrast enhancement. We discuss a number of contrast-enhancement methods and the different angles to approach them, as well as present the properties of the individual methods. The background information presented in this chapter is further supplemented with a discussion of image-based granulometry and oil sand image segmentation, which is discussed in the next chapter.

2.1 Introduction

In this chapter we present background and related work in the area of digital image contrast enhancement. The goal of image contrast enhancement is to improve the quality of an image, so that the image becomes more suitable for a particular application. We find that there is a variety of different approaches to contrast enhancement. We focus on the various concepts that each of these methods explores. Some of these concepts provide the motivation for the contrast enhancement method presented in this thesis. These concepts include multi-scale processing, adaptive filtering and direct enhancement. We first present a general introduction to the area of digital image processing and digital image enhancement.

2.2 Digital Image Processing

Digital image processing is a field that refers to processing of digital images by means of a computer. A greyscale image can be defined as a two-dimensional function $f(x, y)$, where (x, y) are spatial coordinates and f is the intensity at location (x, y) . If x, y and f are finite, discrete quantities, the image is known as *digital image*. A digital image is composed of a finite number of elements, which can be described by x, y and f . Those elements are known as *picture elements* or more commonly as *pixels* [10].

One of the first applications of digital images was in the newspaper industry in the early 1920s. Images were transmitted across the Atlantic via a telegraph wire, where they were

printed at the telegraph receiving station. Digital image processing was not popular until the introduction of digital computers. The first computers powerful enough to be capable of meaningful image processing were developed in the 1960s. It was the development of the computers, and the space program that made digital image processing popular during those years [10].

2.3 Digital Image Enhancement

One major area of digital image processing is image enhancement. Image enhancement deals with improving the quality of images, where the goal is to emphasize wanted features and make them less obscured. Sometimes this is done at an expense of degrading the quality of other details. The area of digital image enhancement is very appealing, where many fundamental image enhancement techniques are built on very simple concepts. A popular example of image enhancement is the enhancement of contrast, since it makes images “look better” [10].

The “look better” part is very subjective. It depends how we define that an image looks better. What is the goal of the enhancement? And how we evaluate that the image is indeed enhanced, and indeed it “looks better”? All that depends on the purpose of the image, and thus the purpose of the enhancement.

The goal of the enhancement can be simply to make the image more appealing to the human eye for an overall pleasing effect, or it could be for making certain details more visible for various applications. One famous area of research dealing with image enhancement is medical imaging. A lot of effort is done to help the radiologists see various details better in medical images so they can provide more accurate diagnosis. Another example of applications of image enhancement is in the area of image segmentation. The goal of the enhancement algorithm in this area is to make it easier to segment individual features of interest in the image.

Whether for human visual or computer image processing applications, image enhancement acts as a pre-processing tool. For human vision, it makes it easier to see individual details. For digital image processing, this pre-processing step makes the subsequent algorithms perform better.

Whatever the application is, the goal of image enhancement is the manipulation of an image in order to improve its quality so that the resultant image is more suitable than the original image for the particular application it is intended for. Image enhancement, as well as de-enhancement, are such preprocessing techniques. Some examples of image enhancement are removing noise, de-blurring, highlighting some specific features needed by a particular application, and of course contrast enhancement [19].

2.4 Contrast Enhancement

What is contrast? In general contrast is defined as a difference in brightness. Often it is the difference between individual objects or features in the image, or between objects and their background. In general, contrast refers to the difference or relationship between the intensity of a given feature and its surroundings. Contrast is what makes objects in an image distinguishable from each other and from the background.

Low-contrast images are often undesirable, since they are more difficult to work with. There can be many causes for low-contrast images. Some examples include: poor illumination, lack of dynamic range in the imaging sensor, or even wrong setting of a lens aperture during image acquisition [10].

Contrast enhancement addresses the problem of improving the contrast in an image in order to make various features more easily perceived. It is used in many applications. One of the most popular applications of contrast enhancement is in the area of medical imaging. Radiologists carefully examine medical images in order to diagnose and treat diseases. Often very subtle features are very crucial in such diagnosis. The goal of contrast enhancement in those applications is to make these subtle features stand out, and thus more easily perceivable by the radiologists.

2.5 Classification of Contrast Enhancement Methods

Not all contrast enhancement methods can be easily classified into distinct categories, since often those methods use aspects from a number of categories. We can assign some attributes to contrast enhancement methods. Contrast enhancement methods can be in spatial domain and in frequency domain. Methods in the spatial domain can be further divided into global and local (adaptive) methods. Another way to classify contrast enhancement algorithms is into direct and indirect. Finally contrast enhancement methods can also be classified as single-scale and multi-scale. Contrast enhancement techniques can be combinations and variations of any of the above.

2.5.1 Global vs. Local Methods

One way to classify spatial contrast enhancement methods is into global and local methods. Global methods process the entire image in the same fashion. On the other hand local (also called adaptive) methods use local image statistics in order to adaptively process a given area in the image. Often image characteristics differ from region to region, so it is reasonable to use context-sensitive approaches when enhancing the contrast. In such approaches the processing adapts to the changing characteristics of the image in the different local regions [19].

2.5.2 Direct vs. Indirect Methods

Contrast enhancement methods are sometimes classified as direct and indirect. The indirect methods do not involve a contrast measure. They do not directly deal with contrast itself, so often the enhanced contrast is a by-product of the operation. The disadvantage of indirect methods is that the contrast (*i.e.* contrast measure) is not directly addressed. There is not much control over the way the contrast is enhanced, or the amount of the enhancement. The enhancement of contrast cannot be controlled directly.

Not being able to control the enhancement of contrast directly may lead to many undesirable effects, such as over-enhancement or under-enhancement. Often these include over-enhancement of noise, over-enhancement of high contrast features, or under-enhancement of low contrast features. For example in the histogram equalization, uniform images or regions are usually over-enhanced, amplifying the noise. Local histogram is discussed in more detail later in this chapter.

Direct methods, as opposed to indirect methods, involve a contrast measure, which is a quantitative measure of contrast in the given image. A contrast enhancement method is called direct when it defines a contrast measure and actually evaluates it during the image enhancement [19].

A contrast measure can be defined globally (*i.e.* a single measure for the entire image) or locally, a separate local measure for each region or even pixel of the image. Local contrast measures are more adequate for images composed of textured regions [2]. The contrast measure yields a value indicating how much contrast there is in a local area (pixel) in the image.

There is no standard way to calculate a contrast measure, and it varies from a method to a method. Usually the contrast measure is highly correlated to the intensity gradient, so it is no surprise that many contrast measures are either defined in terms of intensity gradient or are somewhat similar to gradient operators. In general a local contrast measure is a relationship between the intensity of a foreground feature, such as a pixel, and its contextual background, such a local area surrounding the foreground feature.




Figure 2.1: Many direct contrast enhancement methods use a local contrast measure that deals with computing the relationship between a foreground pixel and its local contextual background.

One example of a contrast measure is the following [6]:

$$c = (l_{max} - l_{min}) / (l_{max} + l_{min}), \quad (2.1)$$

where c denotes the measured contrast, and l_{min} and l_{max} are the minimum and maximum luminance values (respectively) in the area where the contrast is calculated. This area can either be the whole image, or just a part of the image. The values of this measure are in the range $0 \leq c \leq 1$.

A simple example of a direct contrast enhancement method using local statistics is a method first proposed by Wallis and later extended by Lee [19]. This method uses a contrast measure, which is the difference between the grey level G_{xy} of a pixel at location (x, y) , and the grey level mean E_{xy} of local pixels surrounding the pixel at (x, y) . The enhanced grey level G'_{xy} is obtained by multiplying the contrast measure by a factor k as follows:

$$(G'_{xy} - E_{xy}) = k \cdot (G_{xy} - E_{xy}). \quad (2.2)$$

Gordon and Rangayyan [11] define a contrast measure C as follows:

$$C = |p - a| / (p + a). \quad (2.3)$$

In this contrast measure p is the average pixel value in the inner foreground region (whose center is the pixel to be enhanced), and a is the average of all the pixels in the outer background region that surrounds the foreground region. The values of C are in the range $[0, 1]$. The contrast C is enhanced to produce C' using the square root function:

$$C' = \sqrt{C}. \quad (2.4)$$

The square root function ensures that the $[0, 1]$ range for the contrast values is preserved. It also enhances smaller values (lower contrast) more than it does higher values (higher contrast). The value of the enhanced pixel p' is computed as follows:

$$p' = a(1 + C') / (1 - C') \quad \text{if } p \geq a, \quad (2.5)$$

$$p' = a(1 - C') / (1 + C') \quad \text{if } p < a. \quad (2.6)$$

2.6 Different Approaches of Enhancing Contrast

2.6.1 Contrast Stretching


The simplest form of contrast enhancement is known as contrast stretching [10]. The idea behind contrast stretching is to increase the dynamic range of the grey levels in the image being processed. A *transformation function* (also called *mapping function*) is a function that

maps an input grey-level intensity to an output grey-level intensity. Contrast stretching is a transformation function T of the form:

$$s = T(r), \quad (2.7)$$

where r is the input intensity of a pixel, and s is the output intensity of that pixel.

In this method the range of values close to some value m is stretched over the intensity scale, resulting in an improved contrast of those values. This also results in compression of the values that are further away from m . This type of transformation is known as a point operation, since the output value s of any pixel in the image depends only on the input value r of that same pixel. An example of such a mapping function for contrast stretching is shown in Figure 2.2(a). A limiting case of such mapping produces a two-level (binary) image, which is known as thresholding. An example of this form of mapping is shown in Figure 2.2(b).



(a) Typical form of contrast stretching transformation.

(b) A limiting case of contrast stretching, known as *thresholding*.

Figure 2.2: Examples of contrast stretching transformation functions (redrawn from [10]).

2.6.2 Full Frame Histogram Modification

Hall [12], by equalizing the grey-level density of pixels, demonstrated that by modifying the histogram of an image, the perceptiveness of detail can often be increased very significantly. The basic assumption made in histogram modification is that the information contained in an image is related to the probability of occurrence of each grey level. The information content in the image is often easier to perceive, when the probability of occurrence of each grey level is uniformly distributed. A uniform distribution of grey levels tends to make equal use of each quantization level, and to enhance low detail information due to range compression. The histogram equalization technique attempts to obtain this uniform distribution.

Frei [8] suggested the histogram hyperbolization method, which addresses the fact that the response of the human visual system to stimuli is approximately logarithmic. This method of histogram modification redistributes the information in the image hyperbolically.

The full-frame, or global, histogram methods, as pointed out by [22], modify the global histogram of the complete image. These methods are simple to implement, require no interaction from the user, and provide significant visual enhancement of the image. The main problem with these methods is that small, relatively uniform regions may be lost as their grey values are combined with the values of their background. This takes place because the corresponding pixel values, due to their low occurrence, are considered to have low information content. An example of an image enhanced by full-frame histogram equalization is shown in Figure 2.3(c).


2.6.3 Local Area Histogram Modification

The problem of full-frame histogram equalization failing to enhance small local details was first addressed by Ketcham [16], who suggested local-area histogram equalization (LAHE), also known as adaptive histogram equalization (AHE). In this method the enhancement is performed based on a local histogram in a local two-dimensional sliding window. This method was also independently suggested by Pizer [26] for enhancement of medical images, and by Hummel [13]. An example of enhancement using this method is shown in Figure 2.3(e).

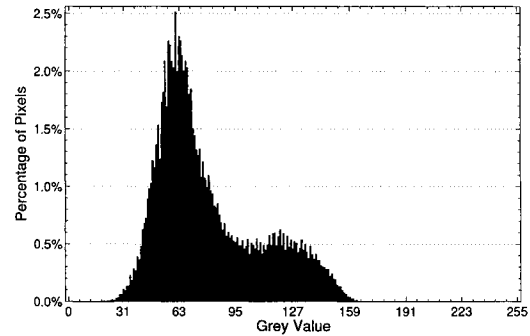
There is a number of problems with the LAHE method. It fails, due to the use of fixed size window, when the window is too small and does not contain multiple objects, or when the window is too large and objects of interest do not occupy a significant portion of the local histogram. Additionally, Rehm and Dallas [29] have shown that LAHE produces undesirable edge artifacts at sharp natural boundaries. This artifact is caused by a change of transformation as the sliding window passes over that boundary. To correct this they suggest to subtract a very smooth version of the image prior to applying the LAHE method, a process called background subtraction.

Contrast-limited adaptive histogram equalization (CLAHE) [25] is a variation of LAHE, where it adds a restriction on how much the contrast is enhanced. This addresses the problem of over-enhancing noise, since there is a limitation placed on the maximum amount of enhancement, thus enhancement of noise as well. This is done by limiting the slope of the mapping function, and in terms of the histogram equalization is equivalent to clipping the height of the histogram. When applying CLAHE the histogram is renormalized, where the clipped portions of the histogram are redistributed evenly through the entire histogram. Enforcing a maximum on counts in the histogram limits the amount of contrast enhancement and thus the enhancement of noise.

Gordon and Rangayyan [11] first suggest contrast enhancement using an adaptive neighbourhood. This algorithm uses a contrast measure which is a ratio between a foreground region centered around a given pixel, and a background region surrounding the foreground



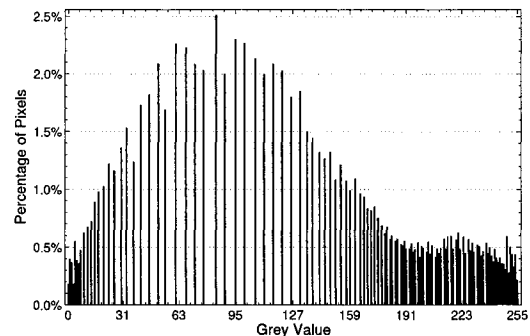
(a) An image of oil sand fragments.




(b) Histogram of (a).




(c) Global histogram equalization of (a).



(d) Histogram of (c).



(e) Local histogram equalization of (d).



(f) Histogram of (e).

Figure 2.3: An example image of crushed oil sand fragments on a conveyor belt and corresponding enhanced images via global histogram equalization and local histogram equalization. The histograms of the three images are also included. Notice how the histogram of the image enhanced by the global method is stretched, whereas the histogram of the image enhanced by the local method has more uniform distribution.

region (refer to Equation 2.3). The idea behind the adaptive neighbourhood is that the size of the neighbourhood is adjusted to be most optimal. This is done by maximizing the contrast measure (*i.e.* which neighbourhood size gives the highest contrast measure). The ratio between the size of the foreground region and the size of the background region is kept constant. Once the size of the adaptive neighbourhood is chosen, the same contrast measure

is used to enhance the contrast of the pixel in the center of the region. Figure 2.4 illustrates an example of the foreground and the background region.




Figure 2.4: Contextual region in Gordon's method (redrawn from [2]).

Paranjape *et al.* [22] propose an adaptive-neighbourhood histogram equalization (ANHE) method. In this method, instead of using a fixed size window, such as in LAHE, an arbitrary shape and size is used for the region from which the histogram is computed. This region is obtained from an adaptive neighbourhood of each pixel being processed, an approach first suggested by Gordon and Rangayyan [11].

In ANHE, an 8-connected set of pixels is grown from the seed pixel, using pixel aggregation, which produces the foreground component of the region from which the histogram is computed. A background component of a constant width is then grown from the outer perimeter of the foreground component, and together with the foreground component constitutes the contextual region of the seed pixel.




Figure 2.5: An example of the contextual region in Adaptive Neighbourhood Contrast Enhancement (ANHE). This region is composed from a foreground component, which includes the seed pixel, and a background component of specified width surrounding the foreground region (redrawn from [19]).

In contrast to LAHE, ANHE is a more adaptive method of contrast enhancement, since it uses a contextually based region, as opposed to a fixed rectangular window used in LAHE. The edge artifacts at sharp natural boundaries, which LAHE tends to produce, are nonexistent, since all the seed pixels in a given region have the same grey-level transformation. ANHE, however, similarly to LAHE, suffers from undesired enhancement of noise, especially in relatively uniform regions.

Adaptive contrast enhancement (ACE) [2] is a method that builds upon Gordon's method

[11], where contrast enhancement is performed over an adaptive neighbourhood. In this method a contrast measure is defined by detecting contours in the image. The measure itself uses local edge detection. The contrast C_{kl} at pixel (k, l) is defined as:

$$C_{kl} = \frac{|X_{kl} - \bar{E}_{kl}|}{|X_{kl} + \bar{E}_{kl}|}, \quad (2.8)$$

where X_{kl} is the value of the pixel, and \bar{E}_{kl} is the mean edge grey level at that pixel location. \bar{E}_{kl} is computed in a local window W_{kl} as follows:

$$\bar{E}_{kl} = \left(\sum_{(i,j) \in W_{kl}} \Delta_{ij} \cdot X_{ij} \right) / \left(\sum_{(i,j) \in W_{kl}} \Delta_{ij} \right), \quad (2.9)$$

where Δ_{ij} is the edge value at pixel (i, j) . The edge value is computed using an edge detection operator, such as the Laplacian:

$$\Delta_{ij} = |X_{ij} - \bar{X}|, \quad (2.10)$$

where \bar{X} is the mean grey level of the eight neighbouring pixels around (i, j) . The contrast is then enhanced using $C'_{kl} = (C_{kl})^{a/b}$, where $b = 2^p$, p is an integer, and $a < b$. Reverse mapping is given by:

$$X'_{kl} = \begin{cases} \bar{E}_{kl} \frac{1 - C'_{kl}}{1 + C'_{kl}} & \text{if } X_{kl} \leq \bar{E}_{kl}, \\ \bar{E}_{kl} \frac{1 + C'_{kl}}{1 - C'_{kl}} & \text{if } X_{kl} > \bar{E}_{kl}. \end{cases} \quad (2.11)$$

This contrast enhancement method is less sensitive to digitization effects and noise. It is useful for extracting contours of objects, where edges are enhanced without significant enhancement of noise.

Mukherjee and Chatterji [19] propose Adaptive Neighbourhood Extended Contrast Enhancement (ANECE). This method combines ANHE [22] and ACE [2]. It performs the same contrast enhancement as ACE, but in contextual regions that are obtained in ANHE. It uses the region-growing approach from ANHE, and the contrast enhancement based on a contrast measure from ACE [2]. The problem of noise over-enhancement is addressed by power variation, where the amount of enhancement is varied between two fixed limits depending on local image statistics. Examples of the basic method, and the power variation version are shown in Figure 2.6. The power variation approach results in a smoother histogram than the basic method.

2.6.4 Unsharp Masking

Unsharp masking is a contrast enhancement technique widely used in photography for over sixty years. It works by making an inverted blurred photographic mask of the original image

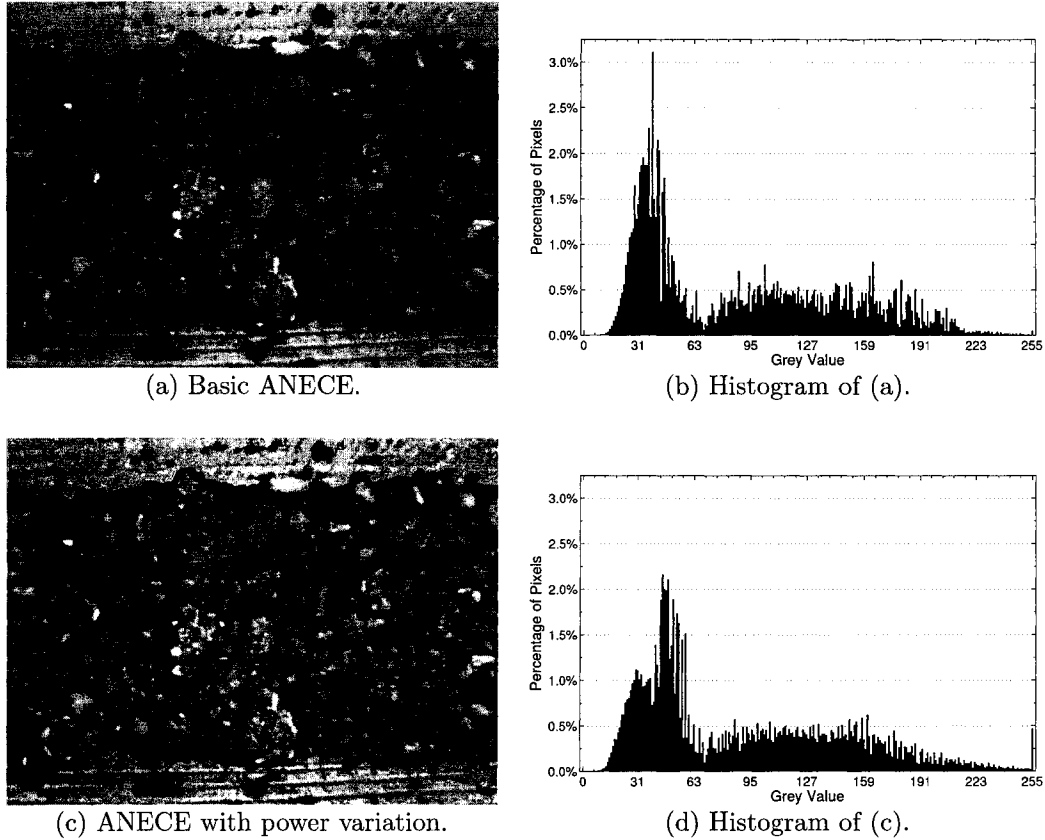


Figure 2.6: An example of the basic ANECE method and its power variation version, enhancing the image from Figure 2.3(a). The histograms of the two images are also included. Notice how the power variation version results in a smoother histogram, decreasing the effect of noise over-enhancement.

on a piece of photographic film. This piece of film is then contact printed in registration with the original image, effectively enhancing edges and small details.

The classic digital unsharp masking is based on adding a high-pass filtered, intensity scaled version of the input image to itself, thus emphasizing the high frequency components of the original image. A *filter* (also called *mask*, *kernel*, *template* or *window*) is a small (*e.g.* 3×3) two-dimensional array of coefficients whose values determine the nature of the operation performed such as blurring or sharpening [10]. The linear unsharp masking filter for the input image $x(n, m)$ and the enhanced image $y(n, m)$ is defined as [28]:

$$y(n, m) = x(n, m) + \lambda z(n, m), \quad (2.12)$$

where $z(n, m)$ is the output of a correction signal obtained via high-pass filtering, and λ is the positive scaling factor that controls the amount of enhancement. The linear unsharp masking suffers from two main drawbacks: it is extremely sensitive to noise, and it enhances high-contrast areas much more than low contrast ones.

A method that addresses these problems is proposed by De Vries [4]. In this method,

sharpening is controlled by an adaptive filter, measuring the input contrast. As a result, the low-contrast areas are enhanced more than the high-contrast ones. This approach still suffers from noise over-enhancement when there is no mismatch between the target and the input dynamic ranges.

Polesel *et al.* [28] introduce an adaptive unsharp masking method. The objective of this method is to emphasize the medium-contrast details more than large-contrast details. Additionally the filter does not perform sharpening operation in smooth areas, making the system more robust to the presence of noise in the input images than traditional approaches. The adaptiveness of the unsharp masking is achieved using a Gauss-Newton adaptation strategy [34] to reduce the squared error between the desired local dynamics and the actual local dynamics.

2.6.5 Contrast Enhancement in Frequency Domain

Tang *et al.* [31] propose a contrast enhancement method using a contrast measure in the Discrete Cosine Transform (DCT) domain. DCT is used in the JPEG image compression format [24]. In this domain the image is divided up into 8×8 blocks, and 64 DCT coefficients are derived for each block, corresponding to different frequencies of the signal within that block. In this contrast enhancement algorithm, those coefficients are grouped into 15 frequency bands. The frequency bands are arranged from lowest to highest frequencies. This algorithm uses a contrast measure c_n defined for each spectral DCT band n , where smaller n corresponds to lower frequencies, greater n corresponds to higher frequencies. This contrast measure is defined as:

$$c_n = \frac{E_n}{\sum_{t=0}^{n-1} E_t}, \quad (2.13)$$

where E_n and E_t are the averages of DCT coefficients in bands n and t , respectively.

The contrast is then enhanced as follows:

$$\bar{E}_n = \lambda H_n E_n, \quad n \geq 1, \quad (2.14)$$


where

$$H_n = \frac{\sum_{t=0}^{n-1} \bar{E}_t}{\sum_{t=0}^{n-1} E_t}, \quad n \geq 1, \quad (2.15)$$

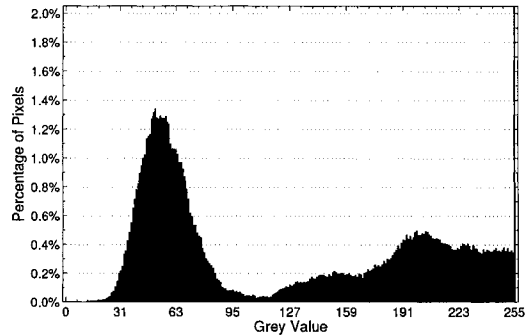
with initial condition $\bar{E}_0 = E_0$. The amount of enhancement is controlled by λ , which is the enhancement control factor with a typical value of 1.95. This spectral form of contrast enhancement gives the algorithm a primitive multi-scale structure.

2.6.6 Multi-Scale Enhancement

In multi-scale contrast enhancement an image is separated (or decomposed) into a number of components, where each component corresponds to a different scale. The actual decomposition method varies, but the choice of it affects how the different features in the image



(a) Enhanced image by non-linear RoLP pyramid recombination.



(b) Histogram of (a).

Figure 2.7: An example of contrast enhancement using non-linear RoLP pyramid recombination, enhancing the image from Figure 2.3(a). Image histogram is also included. This histogram has a bi-modal characteristic. In addition, bright features are overemphasized.

are decomposed. Multi-scale decomposition allows for enhancing of each scale separately, and thus to control the amount of enhancement (or de-enhancement) for each component of the scale.

In direct multi-scale methods, a contrast measure can be defined as a relationship of a foreground feature (*e.g.* a pixel or a small foreground region) to the contextual background region. This contrast measure in a multi-scale representation can be computed using a foreground feature from a finer scale and a background feature from a coarser scale. Furthermore, the multi-scale representation allows for the contrast measure to be computed for different scales, which allows direct enhancement of individual scales. Also, if a multi-scale representation preserves locality information, the contrast measure can be computed locally for individual pixel location, thus allowing an algorithm to be local or adaptive.

Peli and Lim [23] propose an adaptive contrast enhancement algorithm, which separates the image into high and low spatial frequency components, using a low-pass filter. It is an adaptive filter, which modifies contrast locally. The low-pass component determines how much the high-pass component is enhanced. The low-pass component is then modified by a non-linearity, which controls the dynamic range of the image. Both components are then recombined to produce the enhanced image.

Toet [32] proposes an adaptive multi-scale contrast enhancement technique based on recombination of a non-linear ratio of low-pass (RoLP) pyramid. The information present at different levels of the pyramid corresponds to different scales of the original image. This allows for selective enhancement of details at different spatial scales. Each level of the pyramid can be adaptively enhanced using its contextual information from the other levels.

Each level of the RoLP expansion is obtained by applying a low-pass 5×5 Gaussian filter to the image from the previous level, and then downsampling by a factor of 2. The RoLP is

then computed as ratios of the consecutive levels of the expansion. This RoLP representation of the original image is complete, such that the exact image can be reconstructed from the pyramid. The contrast is enhanced during recombination, on all scale-levels, using special recombination rules. Different rules can be used to emphasize different aspects of the encoded information.

Another multi-scale approach to contrast enhancement is in using wavelets. In wavelet-based methods the image is decomposed into a number of discrete frequency bands. Each of these bands can be referred to as scales (thus yielding a multi-scale structure). The advantage of this decomposition is a spatial and frequency representation with multi-scale structure and locality information preserved. Locality allows local/adaptive processing and multi-scale allows scale-based processing. The multi-scale decomposition, if properly chosen, can allow for separation of desired features or objects from noise. In the multi-scale approach each scale can be processed separately allowing for different amount of enhancement or different parameters for each scale.

An example of enhancement using wavelets is a method proposed by Jin *et al.* [15]. In this method over-complete dyadic spline wavelets are used, where the image is decomposed into 2 frequency bands (2 levels) with high frequency and low frequency. A local histogram equalization (LAHE) is performed on each frequency band with a different size of the local histogram window. The advantage of this method comes from combining the local enhancement capability of LAHE, and the selectivity of spatial-frequency components.

A multi-scale morphological approach to local contrast enhancement is proposed by Mukhopadhyay and Chanda [20]. In this method, the scale-space decomposition is obtained via a series of top-hat transformations, varying the size of the structuring element. A top-hat transformation, decomposes an image into a base image and a feature image by applying either morphological opening (white top-hat transformation) or morphological closing (black top-hat transformation). For example the white top-hat transformation is defined as:

$$g(r, c) = \underbrace{(g \circ nB)(r, c)}_{\text{base image}} + \underbrace{[g(r, c) - (g \circ nB)(r, c)]}_{\text{feature image}}, \quad (2.16)$$

where g is the original image, and $(g \circ nB)$ is the morphological opening of g with a disk structuring element B at scale n .

The image is enhanced as follows:

$$\tilde{g}(r, c) = g(r, c) + 0.5 \sum_{i=n}^m F_{iB}^o(r, c) - 0.5 \sum_{i=n}^m F_{iB}^c(r, c) \quad (2.17)$$

for scale range between n and m . $F_{iB}^o(r, c)$ is the white feature image at scale i , and $F_{iB}^c(r, c)$ is the black feature image at scale i . The feature images are defined as follows:

$$F_{iB}^o(r, c) = g(r, c) - (g \circ iB)(r, c), \quad (2.18)$$

$$F_{iB}^c(r, c) = (g \bullet iB)(r, c) - g(r, c). \quad (2.19)$$

In this enhancement technique finer features are emphasized/enhanced more than the coarser features. Greater emphasis of finer features over coarser features makes this method more local and adaptive, since coarser features contribute to more global aspects of this image. In addition, emphasis of fine features results in a sharper image with enhanced edges. An example of enhancement using this method is shown in Figure 2.8. The multi-scale image representation mentioned here is used in our new contrast enhancement algorithm, MMS, and is further discussed in Chapter 4.

This algorithm cannot be applied directly to the enhancement of oil sand images due to a number of limitations. It unconditionally emphasizes finer scales, which leads to over-enhancement of noise and texture. This approach does not offer an ability to control/select the scale and amount of enhancement. Furthermore, it cannot take advantage of a priori size information.




Figure 2.8: An example of contrast enhancement using multi-scale top-hat decomposition, enhancing the image from Figure 2.3(a). The image, due to emphasis of fine features, is sharper, with significant enhancement of texture and noise. Image histogram is also included. The peak at grey value 0 indicates considerable enhancement of black features.

2.7 Summary

In this chapter we present an overview of background and related work in the area of contrast enhancement. Contrast enhancement is used in pre-processing images to improve their quality for specific applications, which can be for the human vision, or for further machine processing.

Contrast enhancement algorithms can be implemented in the frequency domain or spatial domain. Spatial domain algorithms can be global or local (adaptive), where local image statistics are used. Also contrast enhancement algorithms can be direct, which define and evaluate a contrast measure, or indirect. In addition the algorithms can be performed in single-scale or multi-scale image representations.

The classical contrast enhancement methods include contrast stretching, global and local (LAHE) histogram equalization/modification and linear unsharp masking. Methods based on LAHE include CLAHE, which limits the enhanced contrast, and ANHE, which uses adaptive regions of arbitrary shapes. Methods based on unsharp masking include the method by De Vries, which uses an adaptive filter to reduce over-enhancement, and also an adaptive unsharp masking method, which emphasizes medium-contrast details.

Direct contrast enhancement methods define and evaluate a contrast measure during enhancement. Gordon's method uses a contrast measure and variable-size local regions. The adaptive contrast enhancement (ACE) uses the same variable-size local regions, but defines a contrast measure based on local edge detection. The adaptive neighbourhood extended contrast enhancement (ANECE) method combines the contrast enhancement of ACE with the arbitrary-shaped local regions of ANHE.

Other approaches to contrast enhancement include the use of the frequency domain and multi-scale decompositions. The Tang's method enhances an image in the discrete cosine transform (DCT) domain, giving the algorithm a primitive multi-scale structure. Wavelet-based approaches enhance images in multi-scale representations, which also preserve spatial locality. Other multi-scale approaches include the two-scale method by Peli and Lim, non-linear recombination of RoLP pyramid and the method based on multi-scale decomposition using top-hat transformation. Top-hat transformations use morphological operations, which focus on geometric shapes of the features in the image.

As presented in this chapter, there are many approaches to contrast enhancement from many different perspectives, and each having various properties. Some of those properties are used in forming our new contrast enhancement method, MMS, which is discussed in detail in the remaining part of this thesis.

Chapter 3

Image-Based Granulometry of Hard Rocks and Oil Sands

The previous chapter presents an overview of existing contrast enhancement techniques and their various properties. This chapter introduces the problem of oil sand segmentation. It presents some of the existing methods of hard rock and oil sand granulometry, talks about the properties of oil sand images, and points out some of the challenges that are faced with segmentation of those images. The next chapter presents our solution of image contrast enhancement in order to improve oil sand ore image segmentation.

3.1 Introduction

Analyzing oil sand images is very challenging. Oil sand ore comes in a variety of sizes, shapes, colours and textures. Its appearance is greatly affected by the variance in its composition. Crushed oil sand ore comes in a wide range of sizes, where larger objects are often mixed with very fine fragments. Oil sand material is often very soft, resulting in fragment edges not being clearly visible. Varying lighting and weather conditions also play a significant role. The goal of this research is to improve the segmentation of oil sand images through proper contrast enhancement as a necessary first step. Adequate oil sand image segmentation is necessary in providing knowledge to field operators and allowing them to evaluate and optimize oil sand handling equipments. We begin the discussion with a general mentioning of image-based granulometry of oil sand images.

3.2 Image-Based Granulometry

Granulometry is a field that deals mainly with determining the size distribution of particles [10]. The goal of image-based granulometry of oil sand images is to determine the sizes of oil sand particles in an image, and to compute statistical distributions of those sizes.

In contrast to traditional size analysis techniques [7], such as mechanical sieving, cen-

trifugation and sedimentation, it is highly desirable to apply computer vision to oil sand size analysis. Computer image analysis does not interfere with or disrupt the production, and allows analysis of a large number of samples, due to the relatively high speed of image processing. Additionally, a vision-based approach is non-invasive, preserving the shape properties of analyzed oil sand objects.

Maerz *et al.* [18] argue why image-based methods of analysis have many advantages over traditional sieving (screening) methods. Image processing is quick, where many images can be taken and processed quickly. Image processing is also relatively inexpensive in terms of necessary equipment. In addition due to the speed of image processing, lots of samples can be taken and analyzed in a given period of time, reducing the amount of sampling error. Large quantities of material to be analyzed make traditional screening techniques impractical; however image processing is not limited by the volume of the material analyzed. Image processing is non-destructive, whereas in the traditional methods materials can break when screened physically, thus making measurements inaccurate. Image processing also does not interfere with the regular production.

Image-based granulometry of rock materials concentrates mainly on analyzing fragmented hard rocks. Hard rock images are characterized as having well-defined edges and relatively uniform texture. They can be segmented easily using edge-based techniques. An example of an image containing hard rocks is shown in Figure 1.2.




Figure 3.1: An image of oil sand ore on a conveyor belt.

In the case of oil sand images, however, edge-based methods fail due to the rich texture of oil sand fragments, and lack of edge information. In addition, since oil sand mining is a 24-hour, outdoor operation, varying lighting and weather conditions play a significant role in the appearance of oil sand. A size analysis system for oil sand based on segmentation

must resolve a number of technical challenges that are known to be difficult to computer vision. An automated system is desired, where human interaction is not required. Thus such a system needs to be adaptive to the varying visual conditions. An example image of oil sand fragments on a conveyor belt is shown in Figure 1.1 and another example in Figure 3.1.

The typical steps in image-based granulometry are: sampling, digitization, segmentation, size measurement and calculation of size distribution. An image of a scene with material whose particle size distribution is to be measured, is first sampled with an image sensing equipment, typically a CCD camera. Next, the image is digitized, typically with a frame grabber, obtaining a pixel representation of the scene. The digital image is then processed, where individual particles are delineated via image segmentation. The fragment sizes of each segmented particle are measured, and then from these sizes a particle size distribution of the material in the image is calculated. These steps are illustrated in Figure 3.2.

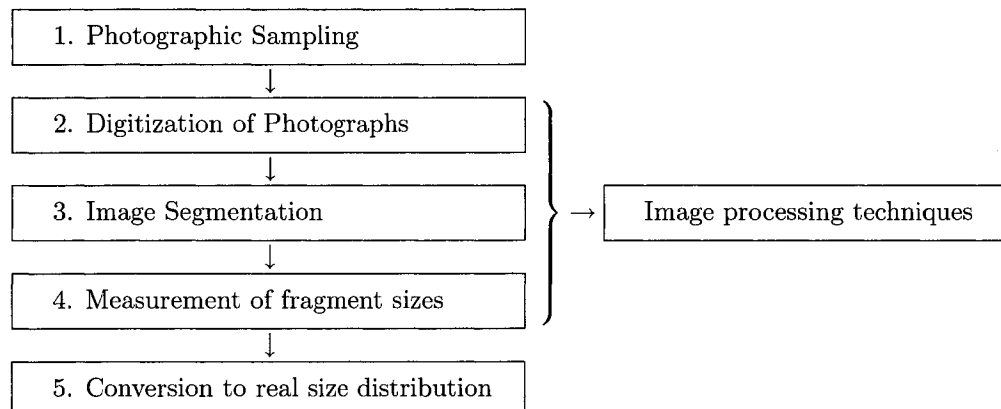


Figure 3.2: Typical steps in image-based granulometry

3.3 Granulometry of Hard Rocks

3.3.1 Measuring Methods

There exists a number of traditional methods (also known as gravitation methods) for measuring the size distribution of blast fragmentation. These include sieving (mechanical and manual), sedimentation and centrifugation. In sieving, coarse material such as crushed rocks, gravels and sands are graded using perforated plate sieves. Fine material like granular soils are graded using woven wire mesh sieves. Centrifugation is a method involving centrifugal force in separating the different mixtures of materials. Sedimentation involves settling of particles in a viscous fluid, where heavier particles settle first, and the reducing density of the suspension is monitored with a floating hydrometer. [7]

3.3.2 Image-Based Measuring Systems

Image-based methods of measuring blast fragmentation have progressed through many stages throughout the history. Initially the analysis was based on photography, where photographs of muckpiles were compared to scaled photographs of standardized muckpiles with known size distribution. Attempts to analyze the size of rock fragmentation from photographs were made as early as the 60's in determining fragment size distribution from photographs of freshly exposed muckpiles and muck surfaces in mine cars [14].

Another example of such an approach is the Compaphoto technique [3], proposed in 1986. This technique assumes that the material distribution follows a Rosin-Rammler curve, given by:

$$R = 100 \exp\left[0.693 \frac{x}{x_m}\right]^n \quad (3.1)$$

where R is the percentage mass retained on screen of size x , x_m is the mean size in the muckpile, and n is the uniformity index. The Rosin-Rammler distribution is an idealized size distribution of material in muckpiles. It is not necessarily true for every case.

Early image-based methods of estimating size distribution of blasted material involve photographic analysis. One such example is a method developed by Anderson [1], which uses grid photography for analyzing the size of fragments from cratering tests with nuclear explosions. In this method a grid is laid out on the muckpile surface and then samples from the muckpile are taken. This method is tedious and involves spending a lot of time on site, thus usually interfering to some degree with the production.

3.3.3 Split[®]

One program developed to analyze size distributions of rock fragments is Split[®] [9]. Split[®] is designed to compute size distributions of greyscale images of hard rock fragments, such as the one shown in Figure 1.2. Split[®] uses algorithms for delineating particles and computing size distributions.

The Split[®] program can be broken into three parts: delineation of the rock fragments, determination of the size and shape of particles with computation of the size distribution, and integration of the two algorithms into two different programs.

The first part of Split[®], which is delineation of the rock fragments, segments a greyscale image into individual particles and the remaining non-particle areas. The typical steps involved are histogram equalization, edge detection, thresholding, followed by object splitting in shadow areas based on edge gradient, and finally a watershed segmentation.

In the second part of Split[®], the size and shape of particles are determined, including computation of size distribution. This is done in a series of steps. First, an actual screen size for each particle is determined. Then this size is corrected for possible overlap and fragment shape using probabilistic methods. Next, the algorithm calculates the overall size

distribution considering all particles in the image. This is followed by approximating the amount of fines not visible due to image resolution, as a percentage of the area of boundary or non-particle area pixels. Finally the algorithm combines the results for a set of images, even if they are captured at different scales.

The third component of Split[®] is the integration of the two algorithms into two separate applications. The first application is an automated system with continuous analysis of material in images of a moving conveyor belt. The purpose of this application is to provide data to operators or to expert systems, such as in a mill or crusher control room.

The second application is a user-friendly off-line program. This program is used to analyze previously captured images, where a user can select different parameters. The primary use of this application is to test various algorithms.

The applications of the Split[®] system include monitoring of the size distribution of particles on a conveyor belt, such as the feed entering a mill, or the product coming out of a primary crusher.

3.3.4 WipFrag[®]

Another edge-based granulometry system used in hard rock industry is WipFrag[®] [18]. WipFrag[®] is an image-based system for analyzing fragment sizes of blasted or crushed rock. It also has applications in granulometry of other materials, such as glass beads or zinc concentrates.

WipFrag[®] analyzes an image and delineates the individual objects or fragments. In this delineation process, WipFrag[®] involves the identification of object edges. This is done in a two-stage process. In the first stage, WipFrag[®] uses a number of common image processing techniques, such as thresholding and gradient operators. This works the best on clean images with lightly textured rock surfaces. In the second stage, WipFrag[®] uses a number of reconstruction techniques to further delineate fragments that are only partially outlined after the first stage. These techniques include knowledge based and arbitrary reconstruction. In addition, WipFrag[®] utilizes the zoom-merge technique [30], which allows to combine results of several images at various scales.

3.4 Granulometry of Oil Sands

Granulometry of Oil Sands has a direct application in oil sand mining. Measurement of oil sand ore size is essential in evaluating and optimizing equipment that handles oil sand feed. Since the nature of oil sand images is significantly different than the nature of images of blasted fragments, generally the approaches used in blast fragmentation do not directly apply to oil sand size measurement.




Figure 3.3: The setup for image-based oil sand granulometry.

There are several challenges in analyzing oil sand images. The ore material comes in a variety of sizes, shapes, colours and textures. Most of the time objects are mixed with dirt and fine particles. Since the images are being captured in an outdoor environment, 24 hours a day, varying lighting and weather conditions play a significant role in the appearance of the objects. The above factors constitute the main challenges of segmenting oil sand objects.

Another challenge with segmenting images of oil sands is the fact that the objects are very close to each other. In many cases they touch one another. With poor edge information, this makes the delineation of particles difficult.

A number of visual properties of the oil sand material in the images can be observed. The apparent brightness of the individual objects varies from object to object, where the majority of larger oil sand chunks tends to appear brighter than the surrounding fines or dirt material, which is the main property that is exploited in the segmentation algorithm [5] used by Ore Size Analyst (OSA), a system designed to measure the size distribution of oil sand particles on a conveyor belt. In addition, the texture of the observed material varies. Portions of the image representing a collection of small objects appear to have a random texture, whereas single large objects appear to have slightly more uniform texture.

3.5 Current Segmentation Method of Oil Sand Images

In this section we discuss the current approach to oil sand image segmentation, and outline a number of problems with which this method is challenged.

3.5.1 Morphological Segmentation

In the past, as [5] points out, edge-based methods have been used in the field of granulometry of fragmented rocks, where edges are clearly visible and textures of individual rock fragments are relatively uniform. In the case of oil sand images, edge-based methods fail due to the complexity of texture and intensity of the analyzed objects. Dornaika and Zhang [5] propose a region-based method, involving mathematical morphology as the means of segmenting object shapes, for use in the segmentation of oil sand images. Morphological segmentation

exploits the spatial properties of the analyzed data and segments the image, delineating individual objects.

3.5.2 Ore Size Analyst

Currently the segmentation of oil sand images is performed using a system developed at the Centre for Intelligent Mining Systems (CIMS), University of Alberta, called the Ore Size Analyst (OSA). OSA is based on original work presented in [5]. The current version of OSA performs a number of image processing steps. First, the region of interest is detected, using an Otsu threshold [21]. This is done in order to separate the background conveyor belt from the foreground oil sand fragments.

The next step is image enhancement. This is the pre-processing stage, that enhances the quality of the image so that it can be better processed by the subsequent steps. In this step the image contrast is enhanced with a local area histogram equalization (LAHE). In this thesis we develop a new algorithm to replace this step, and obtain improved segmentation results.

After enhancement, the image is thresholded by applying a local thresholding technique. The resultant binary image is then segmented, using morphological opening, in order to delineate individual fragments. The segmented image is ran through a number of post-processing steps. One of those steps includes elimination of holes, either at fragment edges or internal to the fragments. Additional optional post-processing steps include further watershed segmentation and elimination of small objects.

There are several problems with the morphological segmentation method, as pointed out in [35]. Frequently the areas of segmented objects are underestimated, especially if the objects are dark and small. Such objects very often end up not being properly segmented. Additional problems include false detection of multiple smaller objects within a larger one, and classification of multiple objects clustered together as a single one. These additional problems are not uncommon.

3.5.3 Contrast Enhancement in OSA

OSA currently uses local area histogram equalization (LAHE) as the contrast enhancement in the image pre-processing step. There are several problems with this method. Although it is an adaptive method, the window size is fixed for the entire image. Regardless of small or large objects in the image, the same window size is used. The method is also indirect (*i.e.* does not use a contrast measure), so there is no direct control over the amount of the enhanced contrast. Another problem with histogram equalization is over-enhancement of relatively low-contrast areas, and over-enhancement of noise. Histogram equalization also enhances internal object texture, which is detrimental to proper segmentation.

Inadequate contrast enhancement contributes to a number of common problems with the segmentation of oil sand images. Those problems include fusion of multiple fragments, and splitting of an individual fragment into smaller pieces. These problems reduce the accuracy and robustness of the oil sand granulometry application. In this thesis we propose a new contrast enhancement algorithm that addresses some of those problems.

3.6 Summary

Image-based granulometry deals with determining size distributions of particles in an image. It has many advantages over traditional (gravitational) techniques of determining particle sizes. Image-based granulometry for analyzing size distributions of rock material mainly concentrates on the hard rock industry, where usually the hard rocks are fragmented by blasting. The properties of such images make it relatively easy to analyze them.

Image-based methods for analyzing rock fragmentation have progressed throughout its history. In the early days there were methods based on photographic sampling; however with the popularity of computers, a number of digital image-based granulometry methods have been developed. Typical processing steps in digital image-based granulometry algorithms are: sampling, digitization, segmentation, size measurement and size distribution calculation.

Two examples of digital image-based granulometry programs for hard rock industry are Split[®] and WipFrag[®]. Split[®] performs image analysis in three stages: delineation of rock fragments, computation of size distribution, and integration into either an automated system, or a manual analysis tool. WipFrag[®] is an edge-based analysis tool for analysis of fragmentation, and performs the analysis in two stages. In the first stage delineation of objects is performed using common image processing algorithms, and in the second stage reconstruction techniques are used to further delineate objects that the first stage did not complete.

The nature of oil sand images differs significantly from the nature of images of hard rocks, and thus methods used in hard rock industry do not directly apply to granulometry of oil sands. Images of oil sands present a lot of challenges in granulometry analysis, which are due to factors such as rich texture of oil sand fragments and poorly defined edges. In addition, oil sand processing is a 24-hour outdoor operation, which is subject to varying lighting and weather conditions.

The current segmentation of oil sand is performed via the Ore Size Analyst (OSA), whose segmentation is based on mathematical morphology. This method takes advantage of properties of oil sand images such as the difference in brightness between larger objects and fine particles or dirt. OSA performs the following steps: region of interest detection, image contrast enhancement, thresholding, morphological segmentation, and removal of holes.

OSA uses LAHE as its contrast enhancement method. LAHE is known to suffer from a number of problems, including noise over-enhancement, inability to enhanced contrast directly and the use of fixed-size local window, which contribute to the common problems with oil sand segmentation by OSA. In this thesis we design a new contrast enhancement algorithm, which addresses some of those limitations.

Chapter 4

Morphological Multi-Scale Contrast Enhancement

The previous chapter discusses image-based granulometry, focusing on problems with the segmentation of oil sand images and the contrast enhancement associated with this segmentation approach. In this chapter we present a new method for contrast enhancement of oil sand images, that addresses some of the problems of the existing OSA method. We outline individual steps performed in this algorithm, and talk about various concepts behind them. In the next chapter we describe the experiments performed to evaluate this method, and provide the results of those experiments.

4.1 Introduction

In this chapter we present the Morphological Multi-Scale (MMS) method of contrast enhancement. This method is based on a morphological multi-scale image decomposition obtained via a series of top-hat transformations. This decomposition allows the choice of scale of enhancement for individual objects, depending on the expected size of each object. The expected object sizes are derived from a priori knowledge, which contains an initial estimation of object segmentation. In addition, the MMS method is direct, where the contrast is measured and controlled during the enhancement. Also, the new method is adaptive, where the enhancement at each pixel location depends on local image statistics at that location. We begin the discussion of this new contrast enhancement method by first presenting some of the basic assumptions we make.

4.2 Basic Assumptions

Our proposed contrast enhancement method MMS is based on several basic assumptions. First of all, objects of different sizes in an image are best enhanced on their corresponding scales. To develop an algorithm capable of this, our proposed method is multi-scale in which

the scale of enhancement is selected based on expected object size. Secondly, only a certain degree of enhancement is required, and both over enhancement and under enhancement are undesirable. To achieve this, we employ a contrast enhancement method that is direct, in which the amount of enhancement is controlled based on a contrast measure. Finally, since objects in different local areas are of different sizes, each local area requires a different degree of enhancement. Our method is spatially adaptive in which different local regions are enhanced differently.

The new contrast enhancement method achieves the enhancement on a current scale. In the scale decomposition usually features corresponding to object contours and features corresponding into internal object textures are separated into different scales. Thus it is advantageous to enhance the objects themselves, without so much enhancing the object texture (rich texture is detrimental to the segmentation).

4.3 Top-Hat Transformation

The MMS algorithm uses top-hat transformations for its multi-scale image representation. Top-hat transformations are tools used to extract either bright or dark features smaller than a given size from an uneven background [20]. Top-hat transformations are defined based on mathematical morphology.

Mathematical morphology is a powerful tool in image processing for extracting image components that represent or describe shape. Morphological operations are based on set theory, treating images as sets of pixels, and performing set operations on the pixels. The two fundamental operations in mathematical morphology are erosion and dilation, and other morphological operations are based on these [10].

In binary images morphological operations of erosion and dilation are defined in terms of logic operations, and in greyscale images in terms of min and max functions. Greyscale erosion $(g \ominus B)(r, c)$ and dilation $(g \oplus B)(r, c)$ are defined as follows [20]:

$$(g \ominus B)(r, c) = \min \{g(r + k, c + l) \mid (k, l) \in B\}, \quad (4.1)$$

$$(g \oplus B)(r, c) = \max \{g(r - k, c - l) \mid (k, l) \in B\}, \quad (4.2)$$

where g denotes the original image and B the structuring element. Also, (r, c) denotes pixel coordinates within image g , and (k, l) denotes pixel coordinates within structuring element B . Morphological opening $(g \circ B)(r, c)$ and morphological closing $(g \bullet B)$ are defined in terms of erosion and dilation as follows:

$$(g \circ B)(r, c) = ((g \ominus B) \oplus B)(r, c), \quad (4.3)$$

$$(g \bullet B)(r, c) = ((g \oplus B) \ominus B)(r, c). \quad (4.4)$$

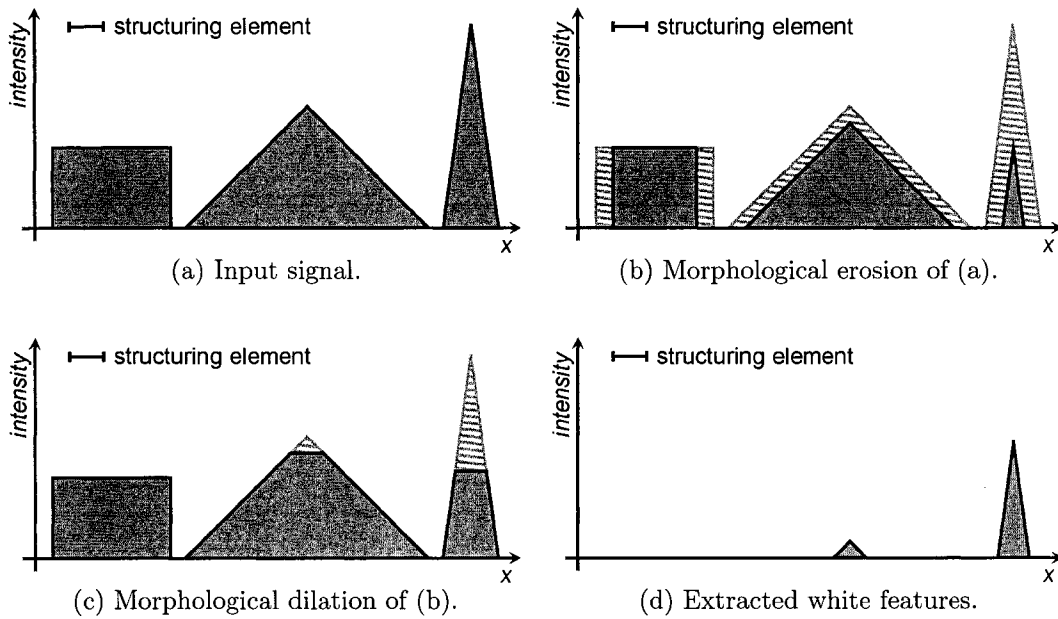


Figure 4.1: An illustration of the greyscale white top-hat transformation.

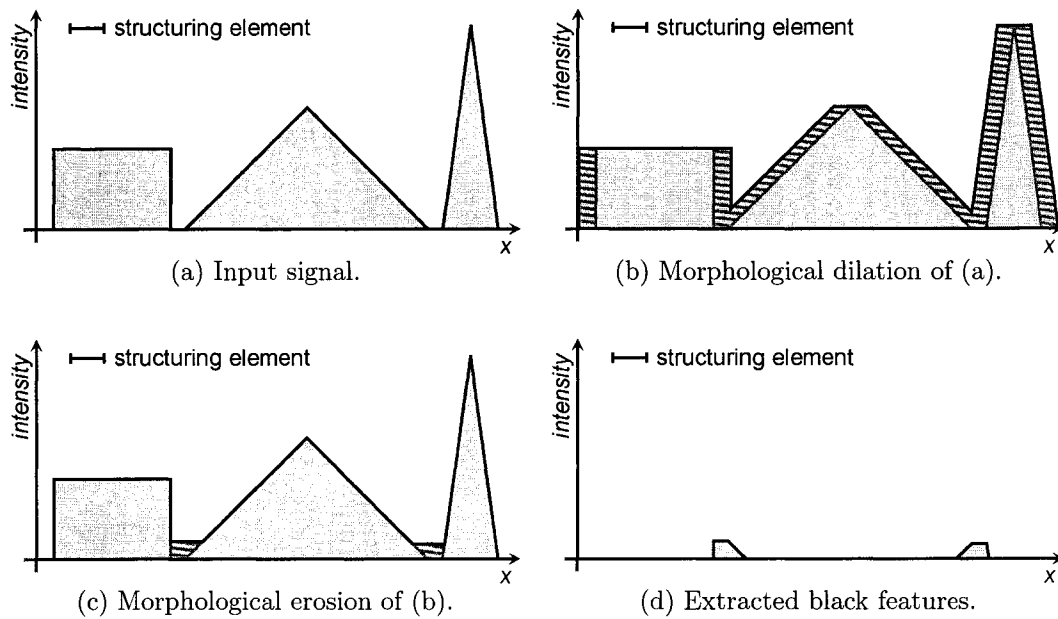


Figure 4.2: An illustration of the greyscale black top-hat transformation (also known as bottom-hat transformation).

A top-hat transformation (as mentioned in Chapter 2) decomposes a given image into a base image and a features image, effectively extracting features that are smaller than a given structuring element. Here, “a feature smaller than the structuring element” refers to a region in which all pixel intensities are greater than those outside the region, and where size as measured by the number of pixels is smaller than that of the structuring element. A white




Figure 4.3: An example of a white and a black top-hat transformation on the image from Figure 3.1.

top-hat transformation which extracts bright features is defined in terms of morphological opening as follows [20]:

$$g(r, c) = \underbrace{(g \circ nB)(r, c)}_{\text{base image}} + \underbrace{[g(r, c) - (g \circ nB)(r, c)]}_{\text{white feature image}}, \quad (4.5)$$

where g is the original image, and nB is a flat disk structuring element B at scale n . The opening operation removes bright features smaller than the structuring element. This is illustrated in Figure 4.1.

While white features are obtained using morphological opening, black features are obtained using morphological closing. This transformation is referred to either as black top-hat transformation or a bottom-hat transformation. It is defined as follows:

$$g(r, c) = \underbrace{(g \bullet nB)(r, c)}_{\text{base image}} - \underbrace{[(g \bullet nB)(r, c) - g(r, c)]}_{\text{black feature image}}, \quad (4.6)$$

Similarly as the opening operation, the closing operation removes dark features smaller than the structuring element. This is illustrated in Figure 4.2. An example of a white and a black top-hat transformation is shown in Figure 4.3.

4.4 Top-Hat Multi-Scale Decomposition

The decomposition part of the algorithm uses top-hat transformations in order to decompose the input image into multiple scales. The scale is controlled by varying the size of the structuring element. The structuring element is a flat disk. The circular shape of the structuring element is chosen since it is invariant to the orientation of objects in the image (isotropic).




Figure 4.4: The top-hat multi-scale decomposition has two independent components, each consisting of decomposed images at different scales. The input image can be exactly reconstructed from any of the two components.

The decomposition is done in two separate (parallel) components: the white features and the black features. Each of these two decomposed components can be fully reconstructed back into the input image (exactly).

In the scale-space decomposition, the individual scales are traversed from lowest (finest) to highest (coarsest). Fine scales correspond to small features; coarse scales correspond to large features. The scale value is represented by the radius of the structuring element disk. The traversing of the scales begins with scale 1, and advances (by 1) to the highest scale (we use 100).

The features are subsequently removed for each scale. The base image is then processed by the next (subsequent scale), where more features are removed, making each feature image a scale slice containing objects greater than or equal to the previous scale and less than the current scale. For example a scale slice obtained at scale x , contains features of the range of $[x - 1, x)$. So features of exactly scale x (and larger) are obtained at the next scale (*i.e.* scale $x + 1$).

The decomposition is partial, that is, if scale range of 1 to 100 is used, objects at scale greater than or equal to 100 are left in the base image. So a scale range from 1 to 100 produces 100 scale slices plus the base image. As a convention, the black feature images are inversed, where higher values represent more black features.

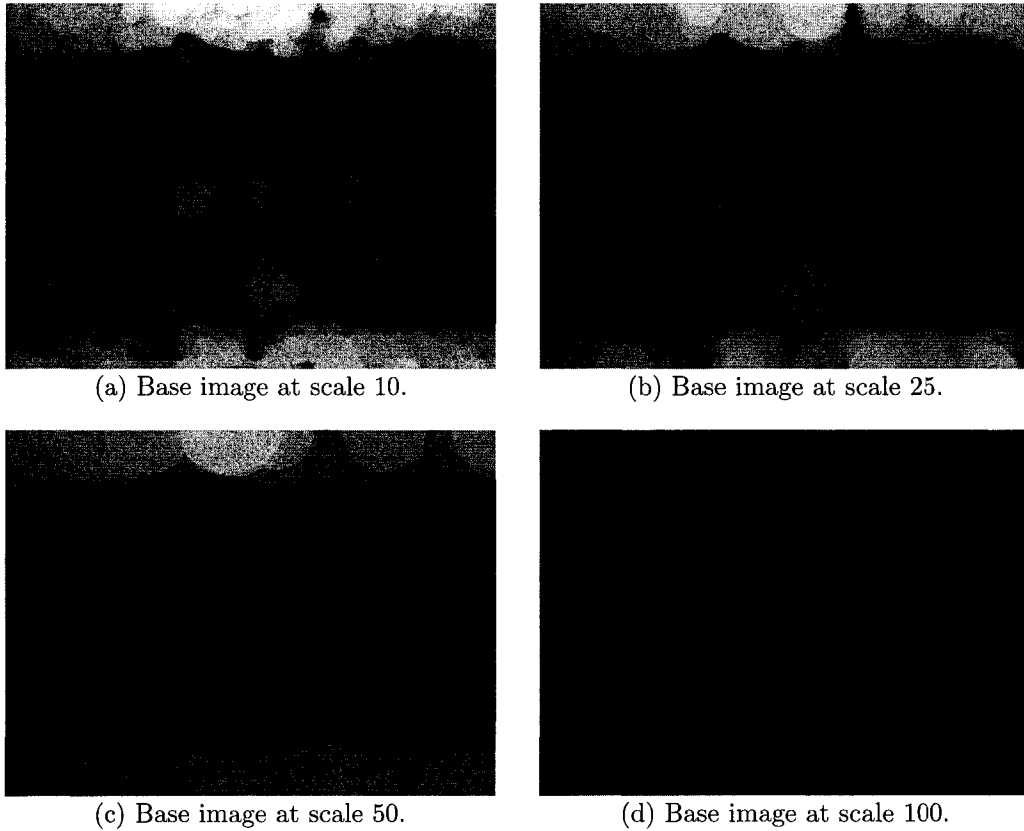



Figure 4.5: An example of base images at different scales of the white-feature top-hat multi-scale decomposition of the image in Figure 3.1. All images are enhanced for viewing purposes.


4.5 Properties of the Multi-Scale Decomposition

The top-hat multi-scale decomposition has several interesting properties. Those properties are illustrated with examples using very primitive objects (*i.e.* flat disks). Considering a number of sample images, artificially generated (as seen in Figure 4.7), containing primitive white objects on black background, we show normalized energy distributions over the first 30 scales of the decomposition.


The most important property of the top-hat multi-scale decomposition (with respect to contrast enhancement of objects of various sizes) is the sensitivity of the multi-scale decomposition to object sizes. This is illustrated in Figure 4.7. In this example only disks of a given size are used, where all objects are of the same size. We obtain normalized energy distribution from the multi-scale decomposition across the first 30 scales. As seen in the example, the peak in the scale distribution corresponds to the radius of the object.




(a) Features slice at scales $5 \leq k < 6$.



(b) Features slice at scales $10 \leq k < 11$.




(c) Features slice at scales $20 \leq k < 21$.




(d) Features slice at scales $40 \leq k < 41$.


Figure 4.6: An example of scale slice images at different scales of the white-feature top-hat multi-scale decomposition of the image in Figure 3.1. All images are enhanced for viewing purposes.




(a) Disk radius $r = 6$.




(b) Energy distribution of (a).



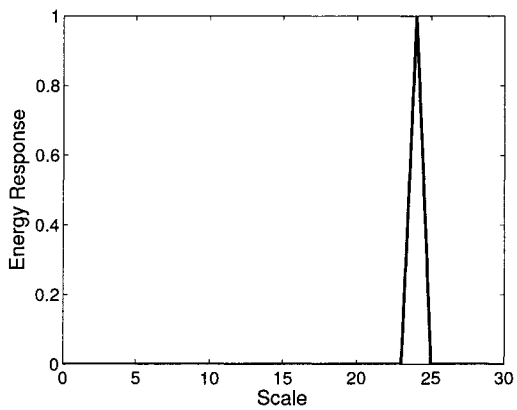
(c) Disk radius $r = 12$.



(d) Energy distribution of (c).




(e) Disk radius $r = 24$.




(f) Energy distribution of (e).


Figure 4.7: Artificial samples with disks of different sizes, and corresponding normalized energy distributions. Notice how the peak in the energy response corresponds with the disk radius.




(a) With Gaussian blur.




(b) Energy distribution of (a).




(c) With noise.



(d) Energy distribution of (c).



(e) With Gaussian blur and noise.



(f) energy distribution of (e)

Figure 4.8: Image from Figure 4.7(c), using disks with radius $r = 12$, with various distortions and corresponding normalized energy distributions. The Gaussian blur is obtained via convolution with a Gaussian kernel with $\sigma = 3$. The noise is generated by adding a random value in the range $[-a, a)$ to each pixel, where $a = 10$.

While being sensitive to object sizes, the multi-scale top-hat decomposition is insensitive to object orientation, number of objects of the same size and shape, object placement and image intensity. The decomposition is insensitive to object orientation, since the shape of the structuring element is a disk. The number of objects of the same size and shape does not affect the normalized energy distribution as well, although it has an effect on the absolute energy distribution. Regardless of the number of such objects, the information corresponding to these objects is distributed in the same way among the individual scale slices. Similarly as the number of objects of the same size and shape, object placement does not affect the distribution of information among the scales. The multi-scale decomposition is also insensitive to image intensity; *i.e.* regardless of the intensity of the image, the distribution of information in the different scales remains the same.

The multi-scale decomposition is partially sensitive to blur, noise and contrast. When objects are blurred, their boundaries become more fuzzy. This increases the spread of information across the different scales. The more blur, the greater the spread, and the energy distribution covers a wider range of scales. Noise in an image corresponds to high-frequency information. Adding noise adds more information to the finest scales.

The decomposition is also partially sensitive to contrast. Higher contrast images tend to have sharper edges. While the peak of the energy distribution still remains around the scale corresponding to the object size, the sharpness of the edge affects the distribution around the peak.

The above properties discussed are for white features only. If the roles are reversed (*i.e.* dark objects on bright background) then the same properties hold for the decomposition of black features.

4.6 Motivation

When we look for a particular object, knowing the object's approximate size and location makes it easier to find it. If somebody tells us that in a particular location there is an object of a certain size (*i.e.* its location and scale), subject to the accuracy of that information, the probability of that location containing objects of the given size rises. Therefore if we know what to expect, we can change the way we look for that object, and make our strategy biased toward those expectations. That way we can detect that object better.

4.7 A Priori Knowledge

The MMS contrast enhancement algorithm uses a priori knowledge in order to determine the best scale of enhancement for each pixel. A priori data is a binary image with delineated regions, where one set of pixels represents the foreground (*i.e.* objects) and the other set the

background. A priori information comes from a previously segmented image, either by OSA or a previous iteration of MMS (as long as the first iteration is not MMS). The previously segmented image acts as an initial state, where the algorithm attempts to enhance the input image, so that in turn the segmentation result can be improved.

Given a priori knowledge (*i.e.* an image with previously segmented regions) MMS produces a *scale map*, which is a matrix of the same dimensions as the input image, that contains scale information. Each value in the matrix indicates the scale of enhancement of the corresponding image pixel. Each segmented region is first labelled with its size, namely the area represented by the number of pixels belonging to that region. Then each pixel area, denoted by a is converted to the radius r of equivalent circle as follows:

$$r = \sqrt{\frac{a}{\pi}}. \quad (4.7)$$

The scale map is then further processed. First, it is blurred using a convolution with a Gaussian kernel. This is done, since the given a priori data is only an approximation of the actual region contours. The Gaussian blurring incorporates the uncertainty into the scale map. The most uncertainty is near the edges of objects, and that is what the Gaussian blurring models. A typical value for the size of the Gaussian kernel, determined experimentally, is $\sigma_p = 2$ pixels.

We perform additional adjustments on the Gaussian blurred scale map, shown in Equation (4.8). Since most objects in the oil sand images are not perfect circles, features corresponding to the object's external contours are usually found on scales a little smaller than the scale corresponding to radius of equivalent circle. In order to compensate for that, we subtract a constant scale reduction value, l , from the scale map. A typical value of l , determined experimentally, is 3.

The final adjustment is to impose a minimum value for each pixel of the scale map. This is done in order to minimize enhancement in the finest scales, which would normally emphasize noise. The typical minimum value, m , is 2, which is determined experimentally. This adjustment has a prominent effect on the background pixels, since the pixels belonging to foreground objects have values which are typically above the minimum.

The adjusted scale map, SM is obtained as follows:

$$SM = \max((SM_G - l), m), \quad (4.8)$$

where SM_G is the Gaussian blurred scale map, l is the scale reduction parameter, and m is the minimum scale map value. Empirically obtained parameters, such as l and m , are summarized in Table 4.1. A pseudo code description of obtaining the scale map is shown in Figure4.11(b).

4.8 Contrast Measure

The MMS algorithm is direct, which means it uses a contrast measure in order to enhance the image. The contrast measure is used to control the contrast itself during the enhancement.

Direct contrast enhancement methods define a contrast measure, and in turn they enhance that measure. The contrast measure is a relationship between a foreground pixel and its surrounding (local) background.

Each input pixel value is transformed into a contrast value, then the contrast value is modified (*i.e.* enhanced), then the modified contrast value is transformed back into a pixel value, now of the enhanced image.

In MMS, the contrast measure c_k in a given scale k is defined as follows:

$$c_k = \frac{I_k}{1 + \sum_{i=k+1}^N I_i}, \quad (4.9)$$

where I_k and I_i are images at scale slices k and i , respectively. N denotes the total number of scales in the decomposition. The corresponding inverse transform is given as:

$$I'_k = c'_k \left(1 + \sum_{i=k+1}^N I'_i \right), \quad (4.10)$$

where I'_k and I'_i are enhanced images at scale slices k and i , respectively. The enhanced contrast measure at scale k is indicated by c'_k .

This definition of contrast measure is based on ideas from [31] (see Equation (2.13)), where in a spectral representation of an image, the higher frequency corresponds the foreground and the lower frequency to the contextual background. As discussed in Chapter 2, typically local contrast measures involve a relationship between foreground features and local contextual background features. In [31] this is achieved by treating higher frequency DCT coefficients as foreground features and lower frequency as background features. The contrast measure is the ratio between these two.

In our contrast measure we treat a given pixel in a given scale slice as foreground, and the same pixel location at coarser scales as the contextual background of the foreground pixel. We use the same idea of the ratio between the foreground and the background (*i.e.* between finer scales and coarser scales). Our multi-scale representation allows this contrast measure to be calculated for any given pixel location. In [31] this locality is limited to individual 8×8 blocks.

In oil sand images, if the scale is matched with the object size, the foreground pixels would belong to features associated with the object being considered, and the background to contextual surroundings of that object. Note also, that the information corresponding to internal object texture would not be present at this scale. Since texture features are smaller than object features, the texture information would be separated into finer scales by the multi-scale decomposition.

The contrast measure tells us how well the foreground pixels stand out from the background pixels. The goal of the contrast enhancement is to increase that ratio, effectively making the object stand out more from its surroundings.

In MMS, the contrast measure is computed for each pixel, for every scale slice image (*i.e.* for each scale). In each scale, the foreground pixel values come from the corresponding scale slice image itself, and the background values come from the summation of all scale slices (already enhanced by now) that are coarser than the current scale. The summation is at the same pixel coordinates as the foreground pixel. The value of the contrast measure is obtained from the ratio between the foreground pixel value and the background of this pixel.

4.9 Control Value

The algorithm uses a *control value* that determines how much to enhance the contrast at a given scale. The control value is calculated from the scale map (discussed in Section 4.7). The scale map (at each pixel coordinate) contains the number of the most appropriate scale to enhance the given pixel at.

Not all scales are enhanced equally, but also not only one scale is enhanced (*i.e.* a peak), because then that causes artifacts. The amount of enhancement (as controlled by the control value) falls off gradually, as we go away from the most appropriate scale. The natural choice to model that fall off is to use a Gaussian function, where the highest peak is at the most appropriate scale. This function is simple enough, and mimics a natural distribution. The Gaussian function, $G(x)$ is defined as:

$$G(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad (4.11)$$

where μ is the mean and σ is the standard deviation of the distribution.

The control value is different for enhancing the white features and the black features. In the enhancement of the white features, we want to focus on enhancing the objects themselves, in the enhancement of black features we want to focus on enhancing the boundaries between objects, but we don't want to emphasize the black features when we are inside the objects, in order to lower the effect of texture (a lowered effect of texture leads to better segmentation). The range of values for the control value is $[0, 1]$, where 0 corresponds to no enhancement and 1 corresponds to maximum enhancement.

In the case of the white features, we enhance mostly on the given scale, representing an object size present in that location. The control value then falls off as we get further away from the most appropriate scale. Also, the coarser the scale, the wider the scale range, and the finer the scale, the narrower the scale range. This is modelled based on the Gaussian





Figure 4.9: The control value function is based on the Gaussian function.



(a) Fine scales.

(b) Medium scales.

(c) Coarse scales.

Figure 4.10: The control value function depends on the expected scale of objects.

distribution. The control value, w_k , for enhancing white features at scale k is given by:

$$w_k = \min \left[\exp \left(\frac{-(s - k)^2}{2(As + B)^2} \right), 1 \right], \quad (4.12)$$

where s is the object scale obtained from the scale map, and A and B are parameters controlling the scale range. The typical values of A and B are determined empirically, and they are $A = 0.06$ and $B = 3.0$.

In the case of the black features, the control value is computed in a similar way as the white features. The control value is reduced the further away from the region scale; however in addition, the value is greatly reduced as we move toward the coarser region scales, and greatly amplified for finer region scales. The range of scales is kept constant (unlike with the white features), but the amplification factor changes (which is constant in the case of white features). The control value for the black features, just like for the white features, is modelled with a Gaussian distribution. This control value, b_k , for enhancing black features at scale k is given by:

$$b_k = \frac{1}{s + 1} \exp \left(\frac{-(s - k)^2}{2C^2} \right), \quad (4.13)$$

where s is the object scale obtained from the scale map, and C is a parameter controlling the scale range. The typical value of C is determined empirically to be 3.0. Empirically obtained parameters are summarized in Table 4.1.

4.10 Image Enhancement

```
INPUT: original image
OUTPUT: two sets {white, black} of decomposed scale slices

load original image
white base image at scale 0 = input image
black base image at scale 0 = input image
for scale k = (1 .. n = 100) {
    create white base image at scale k using Eq. (4.5)
    subtract white base images at scales k-1 and k
    to produce white slice at scale k

    create black base image at scale k using Eq. (4.6)
    subtract black base images at scales k-1 and k
    to produce black slice at scale k
}
```

(a) Top-Hat decomposition.

```
INPUT: a segmented image (a priori)
OUTPUT: scale map

load a segmented image (a priori)
label each region with its size using Eq. (4.7)
apply Gaussian blur with sigma = 2
adjust scale map using Eq. (4.8)
```

(b) Scale map generation.

```
INPUT: two sets of decomposed scale slices, scale map
OUTPUT: MMS contrast enhanced image

for each set of {white, black} scale slices {
    for each pixel position (i, j) {
        lookup scale s at (i, j) in scale map
        for each scale k = (n = 100 .. 1) {
            obtain contrast measure using Eq. (4.9)
            calculate control value using Eq. (4.12) and (4.13)
            enhance the contrast using Eq. (4.14) and (4.15)
            obtain the enhanced pixel value using Eq. (4.10)
        }
    }
}
combine enhanced components {white, black} using Eq. (4.16)
save the image as the MMS contrast enhanced image
```

(c) Image enhancement.

Figure 4.11: Pseudo code of the MMS algorithm.

The MMS algorithm reconstructs the entire decomposition in order to produce the output (enhanced) image. If nothing were to change, then each of the two components of the decomposition (*i.e.* white and black) would be reconstructed into an image identical to the input image. The reconstruction is simply a summation of all the images (*i.e.* scale slices) in a decomposition component.

Image enhancement is performed at every scale during the reconstruction of a component output image. The scales are traversed from the coarsest to the finest (inside a loop with a decreasing index) in order to enhance/reconstruct coarser scales first, before processing finer scales. This ensures the availability of the (reconstructed) background information (*i.e.* scales $> k$) when enhancing any given scale k .

In the case of white features, the contrast measure c_k at a given scale k is enhanced as follows:

$$c'_k = \min [(1 + w_k)c_k + Dw_k, E], \quad (4.14)$$

where c'_k is the enhanced contrast measure. The amount of enhancement is controlled with the control value (for white features) w_k . D and E are parameters that determine how the contrast of white features is enhanced. D determines the sensitivity of the enhancement, and E is the maximum value of the enhanced contrast. The typical values of D and E are determined empirically to be $D = 0.0005$ and $E = 0.0005$.

In the case of the black features the contrast measure c_k at a given scale k is enhanced as follows:

$$c'_k = (1 + Fb_k)c_k, \quad (4.15)$$

where c'_k is the enhanced contrast measure. The amount of enhancement is controlled with the control value (for black features) b_k . The sensitivity of this enhancement is controlled by parameter F . The typical value of this parameter is determined empirically to be 6.0. Empirically obtained parameters are summarized in Table 4.1.

Once the enhanced contrast measure is available, we then convert it back into a corresponding enhanced pixel value, as per Equation 4.10.

We perform this entire procedure for each pixel of the decomposition, for each scale, effectively traversing the entire scale-space decomposition in all 3 dimensions. We do this for both decomposition components (the black and the white decomposition).

After the reconstruction we end up with two enhanced component images. The first enhanced component image, I'_w , comes from the white feature decomposition, and has its white features enhanced. The second enhanced component image, I'_b , comes from the black decomposition, and has its black features enhanced. The final enhanced image, I' , is composed from those two component images as follows:

$$I' = w_w I'_w + w_b I'_b \quad (4.16)$$

Both component images are treated with equal contribution, thus $w_w = w_b = 0.5$. The pseudo code of this algorithm is shown in Figure 4.11.

4.11 Parameters and Their Tuning

The MMS algorithm uses a number of parameters. The values of most of those parameters have been chosen experimentally via a trial and error approach, and are tailored for the particular application of oil sand image contrast enhancement. Table 4.1 lists all the parameters used in this algorithm.

Parameter	Description	Typical value
n_{scales}	Number of scales in the decomposition.	100
σ_p	Scale map Gaussian blurring amount.	2
l	Scale reduction.	3
m	Minimum scale map value.	2
A	Scale range for white features (scale dependent).	0.06
B	Scale range for white features (scale independent).	3.0
C	Scale range for black features.	3.0
D	Object features enhancement level.	0.0005
E	Object maximum contrast.	0.0005
F	Background features enhancement level.	6.0
w_w	Contribution of white features enhancement.	0.5
w_b	Contribution of black features enhancement.	0.5

Table 4.1: Empirically obtained parameters that are used in MMS.

4.12 Fast Computation of the Top-Hat Transformation

In the top-hat multi-scale decomposition the morphological operations use up considerable amount of computing power, especially when dealing with large structuring elements, which we use in the decomposition of coarse scales. In our implementation, we use the Talbot-Van Droogenbroeck method [33] of fast computation of morphological operations. This method implements the basic greyscale morphological operations of erosion and dilation, where an arbitrary shape of the structuring element can be used.

Since erosion is the operation of minimum (and dilation is the operation of maximum) under the structuring element, we track the corresponding minimum or maximum value as we slide the structuring element across the image. As shown in Figure 4.12, if we slide/translate our structuring element for example to the right by 1 pixel, only a very small portion of the structuring element is modified. The pixels that are modified are marked with “-” and “+” symbols. The rest of the values under the structuring element remain unchanged.

In order to track the minimum value (for erosion) or the maximum value (for dilation) the Talbot-Van Droogenbroeck method uses a histogram of the values under the structuring




Figure 4.12: An example illustrating the implementation of shifting of a structuring element. During shifting only edge pixels marked with “+” and “-” need to be considered, and all other pixels are ignored. Pixels marked with “+” are added to a local histogram, and pixels marked with “-” are subtracted.

element, as illustrated in Figure 4.13(a)(b). The minimum grey value is the index of the first non-zero element in the histogram, and the maximum value is the index of the last non-zero element. We compute the initial histogram under the structuring element (*i.e.* for the top left corner of the image). Then all we do is we just slide the structuring element throughout the whole image. Each time we translate the structuring element (by 1 pixel) we only subtract the old values from the histogram, and add the new values to the histogram. As illustrated in the example in Figure 4.12(b), on each translation we subtract the pixel values marked with “-” and add the pixel values marked with “+”.

As an extension to the Talbot-Van Droogenbroeck method, we define a new operation of simultaneous erosion and dilation on the same input image. This is illustrated in Figure 4.13(c). In a single operation we keep track of both the minimum and the maximum value, effectively performing erosion and dilation at the same time.

When we perform the top-hat transform decomposition, we decompose both the white features and the black features, therefore we use both opening and closing on the same input image. The first operation of opening is erosion, and the first operation of closing is dilation. Those first basic operations are both performed on the same input image. In our implementation we combine the erosion of the opening and the dilation of the closing into the new single erosion/dilation operation. This creates a new opening/closing operation, which is implemented using three basic operations, as opposed to four (see Figure 4.14). The new opening/closing takes 75% of the execution time plus a small overhead of the erosion/dilation keeping track of both minimum and maximum values as opposed to only one of them. We find the execution time of our implementation of the opening/closing operation to be 75.85% of the time required to run opening and closing separately.

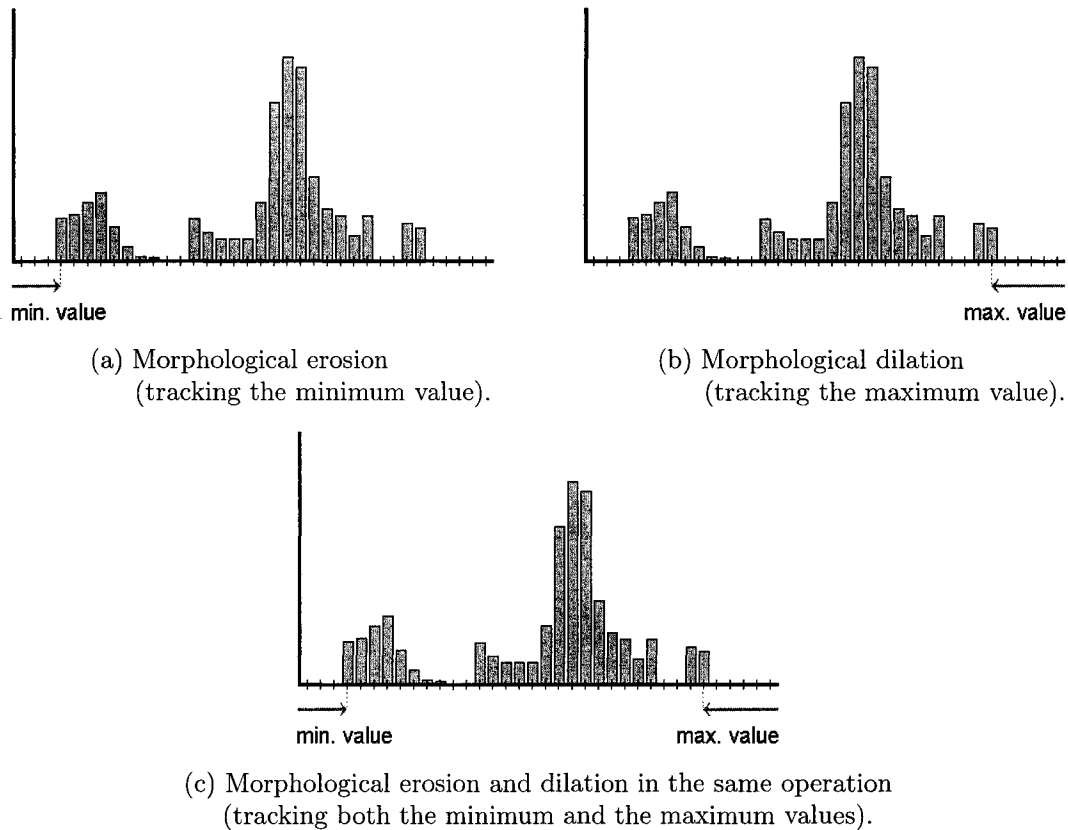


Figure 4.13: An example of a histogram under the structuring element, illustrating the implementation of morphological erosion and dilation in a single operation. In this operation both the minimum and the maximum values are tracked simultaneously.

4.13 Summary

The Morphological Multi-Scale contrast enhancement (MMS) is built on a morphological scale-space decomposition. The scale-space allows for enhancement of individual objects at scales corresponding to their sizes. The multi-scale decomposition is obtained via numerous top-hat transformations, which produce individual scale slices. The decomposition has two components, corresponding to white features and black features. Each of those two components of the decomposition can be independently reconstructed back into the original image.

The white features decomposition component is obtained via white top-hat transforms, which is based on morphological opening. Similarly the black features component is obtained via black top-hat transformation, based on the closing operation.

The multi-scale top-hat decomposition has a number of interesting properties. Most importantly, the decomposition is sensitive to object size, which allows the choice of scale of enhancement for each object, depending on its size. Also, the decomposition is insensitive to object orientation, object count, object placement and intensity, and is partially sensitive

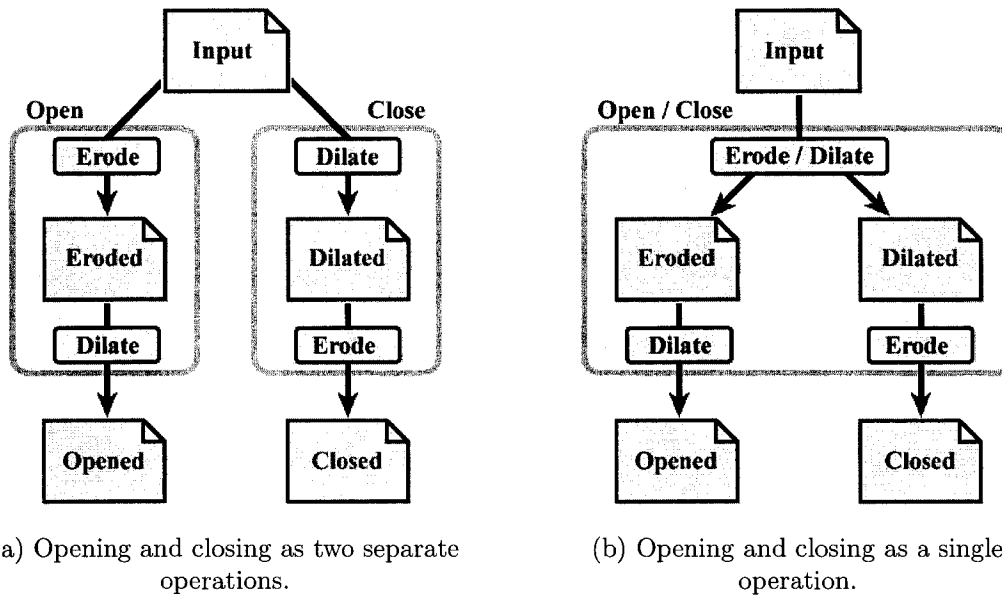


Figure 4.14: An example illustrating the implementation of morphological opening and closing in a single operation. The erosion in the opening and the dilation in the closing both take the same image as input, thus they can be combined into a single erosion/dilation operation.

to blur, noise and contrast.

The MMS algorithm uses a priori knowledge as the initial information about the objects in the image. This information is in the form of a segmented binary image. This image is then converted into a scale map, where each pixel is assigned the scale value corresponding to the object it belongs to, thus effectively selecting the scale of enhancement.

This algorithm is direct. It uses a contrast measure to perform the enhancement. The measure is computed from the input decomposition, enhanced, and then converted back into pixel values of the enhanced multi-scale decomposition.

The amount of enhancement at each scale is specified by a control value, which is computed from the scale map and the enhancement scale. It assures that the scales corresponding to the object scales at each pixel location are enhanced the most, and the neighbouring scales are also enhanced, but to a lesser extent. The amount of enhancement in the neighbouring scales depends on how much the object scales and the enhancement scales differ, and is modelled with a Gaussian distribution.

During the enhancement the entire scale-space decomposition is processed at every scale. Both white and black feature components are enhanced, and combined together to produce the output image. The enhancement parameters are chosen on a trial and error basis, and are tuned to the particular application of oil sand image contrast enhancement.

The top-hat transformations are implemented using an optimized method of computing erosion and dilation, which involves a local histogram in a sliding window of the shape

of the structuring element. Our additional optimization takes advantage of the fact that when erosion and dilation are both computed from the same input image, they can be combined into a single operation. This reduces the total time of obtaining the multi-scale decomposition by almost a quarter.

Chapter 5

Experiments and Results

The previous chapter talks about the new MMS method of enhancing contrast of oil sand images. In this chapter we describe the experiments performed to evaluate this method. We talk about ground-truth images, scoring metric and experimental setup. We also present the results of the experiments. The next chapter presents the conclusion of this theses, including the lessons learned and suggestions for future work.

5.1 Introduction

Since the goal of the contrast enhancement is the improvement of the segmentation performance, the MMS algorithm is evaluated in terms of the segmentation performance. Experiments are set up to compare the segmentation while using the new MMS algorithm for contrast enhancement to the segmentation while using the current contrast enhancement method of OSA (which is local histogram equalization). We also present experimental results, which illustrate the improvement of oil sand image segmentation when using MMS to enhance image contrast. We show even further improvement, when the MMS algorithm is performed in more than one iterations. The improvement is measured quantitatively with a scoring metric. The scoring metric evaluates, on a per-object basis, how close a segmented image is to a provided manually generated ground-truth image. We first talk about the ground-truth images, which are essential in evaluating the new contrast enhancement algorithm.

5.2 Ground-Truth Images

In order to evaluate the segmentation performance of the system, 10 manually-segmented images are used. These images have been carefully segmented by hand, and are used as ground-truth images for the comparisons. Each of the ground-truth images is binary, having only labels of 0 (black) and 1 (white). Pixels with label 0 are background pixels, and pixels with label 1 are foreground pixels (*i.e.* belonging to the segmented regions). All segmented




Figure 5.1: An example of an oil sand input image and its corresponding manual ground-truth segmentation.

regions (composed of foreground pixels) are delineated, *i.e.* have at least 1 pixel (composed of background pixels) separation between them. Figure 5.1 shows an example of an input image and a corresponding segmented ground-truth image. What is measured in the experiments is how close the segmentation results get to this ground truth.

5.3 Evaluation using Scoring Metric

While most contrast enhancement techniques are evaluated qualitatively, we attempt to evaluate our technique quantitatively. Typically the evaluation is tied to the purpose, goal or application of the method in question. Our contrast enhancement method is designed for oil sand segmentation applications, and therefore is evaluated in terms of segmentation performance.

The segmentation performance in the experiments is evaluated in terms of the scoring metric [27]. The scoring metric compares two binary (segmented) images, and gives a single score value that says how well the two images match. One of the most basic ways to compare how well two binary images match is a intersection over union ratio:

$$IU(A, B) = \frac{|A \cap B|}{|A \cup B|}, \quad (5.1)$$

where A and B are the two binary images being compared. The intersection of A and B (*i.e.* $A \cap B$) is a set of only those foreground pixels found in both of these two images. The union of A and B (*i.e.* $A \cup B$) is a set of foreground pixels found in any (or both) of the two images. $|A \cap B|$ denotes the number of pixels in the intersection of A and B , and $|A \cup B|$ denotes the number of pixels in the union of A and B . The intersection over union ratio yields a value in the range $[0, 1]$.

Intersection over union is a pixel-based method. When applied to entire images it does not evaluate the match in terms of objects or regions, which is what segmentation produces.

The scoring metric we use performs object-based comparison of two segmented images, considering individual segmented regions. Instead of applying intersection over union to entire images, this score metric applies intersection over union to individual object pairs, where each object in the pair comes from the other image being compared. The score metric $S(A, B)$ between two binary images A and B is defined as follows [27]:

$$\delta(x) = \begin{cases} 0 & \text{if } x = 0 \\ 1 & \text{if } x \neq 0 \end{cases} \quad (5.2)$$

$$S_{\text{partial}}(A, B) = \sum_{i=1}^m \left[\sum_{j=1}^n \frac{|A_i \cap B_j|}{|A_i \cup B_j|} \times \frac{\delta(|A_i \cap B_j|)|B_j|}{\sum_{k=1}^n \delta(|A_i \cap B_k|)|B_k|} \right] \frac{|A_i|}{\sum_{k=1}^m |A_k|} \quad (5.3)$$

$$S(A, B) = \min(S_{\text{partial}}(A, B), S_{\text{partial}}(B, A)) \quad (5.4)$$

where m is the number of fragments in A and n is the number of fragments in B . An input binary image A having m fragments is defined as:

$$A = \{A_1, A_2, \dots, A_m\}, \quad (5.5)$$

where A_i is the i^{th} fragment in A . The number of foreground pixels in fragment A_i is denoted by $|A_i|$. The definition of image B is analogous to image A .

The segmentation score between two images A and B is calculated by traversing each segmented region, A_i , of image A , and computing the intersection over union ratio between A_i and each region, B_j , in image B that overlaps A_i by at least 1 pixel (*i.e.* $|A_i \cap B_j| \neq 0$). The total score for the region A_i is calculated by a weighted sum of all the scores for overlapping regions. The weights are determined by the area (represented in number of pixels) of the overlapping regions. This region score is calculated for every segmented region A_i , and the final score is a weighted sum of all those region scores. These weights are determined by the area of the corresponding regions in image A .

The scoring metric has a number of properties. It produces a single value in the range $0 \leq S(A, B) \leq 1$, that indicates the quality of the segmentation, or namely the similarity between two segmented images A and B . Score metric value of 1 indicates perfect match, that is:

$$S(A, B) = 1 \quad \text{if and only if } A = B. \quad (5.6)$$

Value of 0 indicates no match at all, that is:

$$S(A, B) = 0 \quad \text{if and only if } A \cap B = \emptyset \text{ and } A \cup B \neq \emptyset. \quad (5.7)$$

Additional properties of the score metric include the fact that in the metric error contribution of an object is proportional to its size. Also, the scoring metric is symmetric, that is:

$$S(A, B) = S(B, A). \quad (5.8)$$

The metric is also scale invariant, that is:

$$S(A, B) = S(nA, nB), \quad (5.9)$$

where nA represents an image A scaled by a factor $n > 0$.

The score value is always higher for a better match and lower for a worse match, which holds for the following conditions: over-segmentation, under-segmentation, misalignment, over-size, under-size, spurious objects and missing objects. These conditions are listed in Table 5.1.

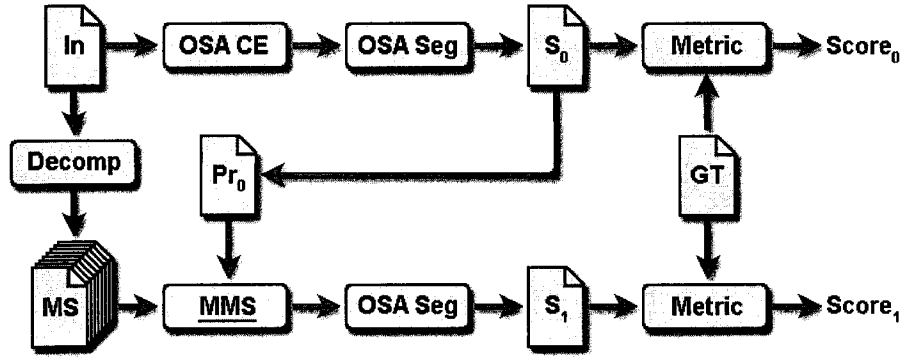
Condition	Description
over-segmentation	a single object incorrectly split into two or more smaller fragments
under-segmentation	several objects incorrectly merged into one larger fragment
misalignment	segmented pieces in incorrect locations
over-size	segmented pieces too large
under-size	segmented pieces too small
spurious objects	incorrect introduction of non-existing fragments
missing objects	valid fragments not included in the segmentation

Table 5.1: Segmentation conditions handled by the score metric.


5.4 Experimental Setup

The experiments are designed to compare the segmentation performance using the original OSA vs. using the new contrast enhancement. The segmentation performance is measured in terms of the score metric.

The initial run comes from the original OSA algorithm. The segmented output serves as a priori knowledge to the new contrast enhancement algorithm. This is the first iteration of the new contrast enhancement algorithm. The output of this segmentation's run is used as a priori input to the next iteration, and so on. This experiment is carried out for several iterations. The whole procedure is repeated for all 10 images.



(a) Experimental setup for the first iteration of MMS.



(b) Experimental setup for subsequent iterations of MMS.

Figure 5.2: Experimental setup for testing of the MMS algorithm. Abbreviations used in this figure are explained in Table 5.2.


Abbreviation	Description
Decomp	multi-scale top-hat decomposition
GT	ground-truth image
In	input image
Metric	segmentation score metric
MMS	the MMS contrast enhancement algorithm
MS	multi-scale decomposed slices
OSA CE	OSA's contrast enhancement
OSA Seg	OSA's segmentation
P_0	a priori data from the initial OSA's iteration
P_{i-1}	a priori data from the $i - 1$ iteration of MMS
S_0	segmented image from the initial OSA's iteration
S_1, S_i	segmented image from 1 st and i^{th} iteration of MMS, respectively
$Score_0$	segmentation score from the initial OSA's iteration
$Score_1, Score_i$	segmentation score from 1 st and i^{th} iteration of MMS, respectively

Table 5.2: Explanations of abbreviations used in Figure 5.2.

5.5 Results




(a) Enhanced image by local histogram equalization.



(b) Segmentation of (a) - score: 0.5594.




(c) First iteration of enhancement by MMS.



(d) Segmentation of (c) - score: 0.6443.




(e) Sixth iteration of enhancement by MMS.




(f) Segmentation of (e) - score: 0.6707.

Figure 5.3: Example results for Image A, from Figure 5.1(a). This examples includes the initial iteration of OSA, and the first and sixth iteration of MMS. Corresponding segmented images are also shown. A complete set of result images for all 10 input images is included in Appendix A.




(a) Enhanced image by local histogram equalization.



(b) Histogram of (a).




(c) First iteration of enhancement by MMS.



(d) Histogram of (c).



(e) Sixth iteration of enhancement by MMS.



(f) Histogram of (e).

Figure 5.4: Example results for Image A, from Figure 5.1(a). This examples includes the initial iteration of OSA, and the first and sixth iteration of MMS. Corresponding image histograms are also shown. The histograms from the MMS method have a bi-modal characteristic, showing the separation between the bright and the dark features. Subsequent iterations of MMS result in a smoother histogram. MMS enhances some pixels to the maximum grey value, which can be seen by the peaks at grey value 255 in the corresponding histograms.

Image	OSA	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$	$i = 6$	GT
Image A	329	219	235	229	227	229	227	213
Image B	281	187	193	195	197	189	201	171
Image C	347	233	233	239	233	235	237	237
Image D	307	225	237	237	237	229	235	213
Image E	347	275	279	277	285	273	277	303
Image F	317	211	207	203	195	201	191	171
Image G	315	215	219	215	213	211	205	207
Image H	337	199	223	217	219	213	207	213
Image I	363	273	255	269	261	275	261	273
Image J	375	289	309	307	297	297	299	309
Mean	331.8	232.6	239.0	238.8	236.4	235.2	234.0	231.0
Variance	781.5	1204.3	1178.7	1262.6	1206.3	1254.6	1263.3	2424.0
95% conf.	17.3	21.5	21.3	22.0	21.5	22.0	22.0	30.5

Table 5.3: Numbers of segmented fragments for the first 6 iterations ($1 \leq i \leq 6$) of MMS. The numbers of fragments of the initial OSA iteration, as well as using the ground-truth as a priori data are also included.

Image	OSA	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$	$i = 6$	GT
Image A	0.5594	0.6443	0.6461	0.6590	0.6648	0.6667	0.6707	0.7854
Image B	0.6413	0.6487	0.6502	0.6636	0.6531	0.6580	0.6497	0.7961
Image C	0.5873	0.6476	0.6577	0.6667	0.6563	0.6563	0.6468	0.7687
Image D	0.4952	0.5341	0.5369	0.5505	0.5559	0.5605	0.5631	0.7284
Image E	0.5909	0.6347	0.6528	0.6331	0.6565	0.6486	0.6509	0.7449
Image F	0.5877	0.6259	0.6512	0.6553	0.6497	0.6347	0.6365	0.8243
Image G	0.5159	0.5586	0.5662	0.5734	0.5743	0.5777	0.5795	0.7978
Image H	0.5810	0.6467	0.6372	0.6478	0.6459	0.6562	0.6592	0.7612
Image I	0.5430	0.6109	0.6291	0.6272	0.6394	0.6362	0.6406	0.7590
Image J	0.5543	0.6223	0.6249	0.6261	0.6265	0.6289	0.6377	0.7641
Mean	0.5656	0.6174	0.6252	0.6303	0.6322	0.6324	0.6335	0.7730
Variance	0.0017	0.0016	0.0017	0.0015	0.0014	0.0013	0.0012	0.0008
95% conf.	0.0259	0.0247	0.0253	0.0243	0.0230	0.0221	0.0214	0.0175

Table 5.4: Results of the experiment for the first 6 iterations ($1 \leq i \leq 6$) of MMS. The results of the initial OSA iteration, as well as using the ground-truth as a priori data are also included.

The results from running the experiments are obtained. An example set of results, for Image A (from Figure 5.1(a)), is shown in Figure 5.3. Figure 5.3(a) shows the output of contrast enhancement of the initial iteration of OSA, which is local histogram equalization. The corresponding segmentation is shown in Figure 5.3(b). The segmentation score for that particular image is 0.5594.

Figures 5.3(c)-(f) show example results, for the same image, for the first and sixth iterations of MMS. The segmentation score for this image after the first iteration is 0.6443, and after the sixth is 0.6707. The sixth iteration of MMS produces the maximum segmentation score for this image. The complete set of resultant images for all 10 input images is included in Appendix A.

Image	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$	$i = 6$	GT
Image A	0.0849	0.0867	0.0996	0.1054	0.1073	0.1113	0.2260
Image B	0.0074	0.0089	0.0223	0.0118	0.0167	0.0084	0.1548
Image C	0.0603	0.0704	0.0794	0.0690	0.0690	0.0595	0.1814
Image D	0.0389	0.0417	0.0553	0.0607	0.0653	0.0679	0.2332
Image E	0.0438	0.0619	0.0422	0.0656	0.0577	0.0600	0.1540
Image F	0.0382	0.0635	0.0676	0.0620	0.0470	0.0488	0.2366
Image G	0.0427	0.0503	0.0575	0.0584	0.0618	0.0636	0.2819
Image H	0.0657	0.0562	0.0668	0.0649	0.0752	0.0782	0.1802
Image I	0.0679	0.0861	0.0842	0.0964	0.0932	0.0976	0.2160
Image J	0.0680	0.0706	0.0718	0.0722	0.0746	0.0834	0.2098
Mean	0.0518	0.0596	0.0647	0.0666	0.0668	0.0679	0.2074
Variance	0.0005	0.0005	0.0005	0.0006	0.0006	0.0008	0.0016
95% conf.	0.0136	0.0141	0.0136	0.0154	0.0153	0.0175	0.0249

Table 5.5: Absolute improvement, derived from Table 5.4.

Image	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$	$i = 6$	GT
Image A	0.1518	0.1550	0.1780	0.1884	0.1918	0.1990	0.4040
Image B	0.0115	0.0139	0.0348	0.0184	0.0260	0.0131	0.2414
Image C	0.1027	0.1199	0.1352	0.1175	0.1175	0.1013	0.3089
Image D	0.0786	0.0842	0.1117	0.1226	0.1319	0.1371	0.4709
Image E	0.0741	0.1048	0.0714	0.1110	0.0976	0.1015	0.2606
Image F	0.0650	0.1080	0.1150	0.1055	0.0800	0.0830	0.4026
Image G	0.0828	0.0975	0.1115	0.1132	0.1198	0.1233	0.5464
Image H	0.1131	0.0967	0.1150	0.1117	0.1294	0.1346	0.3102
Image I	0.1250	0.1586	0.1551	0.1775	0.1716	0.1797	0.3978
Image J	0.1227	0.1274	0.1295	0.1303	0.1346	0.1505	0.3785
Mean	0.0927	0.1066	0.1157	0.1196	0.1200	0.1223	0.3721
Variance	0.0015	0.0017	0.0016	0.0021	0.0021	0.0027	0.0089
95% conf.	0.0244	0.0252	0.0249	0.0284	0.0286	0.0325	0.0585

Table 5.6: Relative improvement, derived from Table 5.4.

Image	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$	$i = 6$
Image A	0.3757	0.3836	0.4407	0.4664	0.4748	0.4925
Image B	0.0478	0.0575	0.1441	0.0762	0.1079	0.0543
Image C	0.3324	0.3881	0.4377	0.3804	0.3804	0.3280
Image D	0.1668	0.1788	0.2371	0.2603	0.2800	0.2912
Image E	0.2844	0.4019	0.2740	0.4260	0.3747	0.3896
Image F	0.1615	0.2684	0.2857	0.2620	0.1986	0.2063
Image G	0.1515	0.1784	0.2040	0.2072	0.2192	0.2256
Image H	0.3646	0.3119	0.3707	0.3602	0.4173	0.4340
Image I	0.3144	0.3986	0.3898	0.4463	0.4315	0.4519
Image J	0.3241	0.3365	0.3422	0.3441	0.3556	0.3975
Mean	0.2523	0.2904	0.3126	0.3229	0.3240	0.3271
Variance	0.0124	0.0139	0.0100	0.0148	0.0139	0.0182
95% conf.	0.0691	0.0730	0.0621	0.0755	0.0731	0.0836

Table 5.7: Potential captured, derived from Table 5.4.

Table 5.3 contains the number of fragments for the segmented images of the initial OSA iteration, first 6 iterations of MMS and the ground-truth images corresponding to each of the ten sample images. On average, there are 331.8 fragments in the initial segmentation by OSA, with variance of 781.5. The number of fragments is less for the first 6 iterations of MMS, ranging from 232.6 to 239.0 with variance ranging from 1178.7 to 1263.3. The ground-truth images have on average 231.0 fragments with variance of 2424.0.

The experimental results, in terms of segmentation scores, for all 10 images are shown in Table 5.4. The average score for the initial iteration of OSA is 0.5656 with variance of 0.0017. Using the output of the initial segmentation as a priori data to the MMS algorithm, the average score for the first iteration of the algorithm is 0.6174 with variance of 0.0016. This produces an absolute average improvement of 0.0518, which is a relative improvement of 0.0927.

The maximum average score (for the 10 images in question) is produced by the sixth iteration of MMS, and it is 0.6335. This is an absolute improvement of 0.0679, or a relative improvement of 0.1223. Table 5.4 lists the results for the first 6 iterations of MMS. In addition Tables 5.5 and 5.6, which are derived from Table 5.4, list the absolute and the relative improvement, respectively, of each iteration of MMS. The absolute and relative improvements are calculated with respect to the score obtained via the initial iteration of the OSA algorithm.

The tables mentioned above contain an additional score column for using ground-truth images (denoted as “GT”) as a priori data. In this case, MMS uses the manually generated ground-truth segmentation image directly as a priori data input. This column denotes the maximum potential of this algorithm when “perfect” a priori data is given.


Based on the experimental results, we calculate the *potential captured* for each image, which is a measure of the algorithm’s performance with respect to the given a priori data. In this measure, we assume that the minimum potential of 0 is given by the score of the initial iteration of OSA, and the maximum potential of 1 is given by the score when the ground truth image is supplied as a priori data. The potential captured p_c for a given score S_X is calculated as follows:

$$p_c = \frac{S_X - S_{OSA}}{S_{GT} - S_{OSA}}, \quad (5.10)$$


where S_{OSA} denotes score from the initial iteration of OSA, and S_{GT} denotes the score given the corresponding ground-truth image as a priori input. This potential captured is shown in Table 5.7, where average potential captured after the first iteration is 0.2523, and after the sixth iteration is 0.3271.

The experiments have been performed on each of the images for the first 50 iterations of MMS. A plot of the average score for all 10 images is presented in Figure 5.5(a). In this plot, the initial iteration of OSA is denoted by 0, and the subsequent iterations (1-50)

denote the iterations of MMS. The first iteration of MMS produces the most improvement compared to the initial iteration of OSA. The maximum average score occurs at iteration 6. From there the average score decays, but it does establish a relatively stable level.



(a) Average segmentation score for the first 50 iterations of MMS.



(b) Average consecutive matching score for the first 50 iterations of MMS.

Figure 5.5: Results for the first 50 iterations of average score for Images A-J. The most improvement is produced by the first iteration of OSA. The maximum average segmentation score of 0.6335 is observed in iteration 6 of MMS. In the subsequent iterations, the average segmentation score is lower, but remains relatively uniform. Individual plots for each of the 10 test images are included in Appendix A.

Figure 5.5(b) shows the average consecutive matching score for the 10 test images for the first 50 iterations. The consecutive matching score is calculated using the score metric described in Section 5.3; however instead of comparing a given segmentation image against the ground-truth, the segmentation for a given iteration of MMS is compared against the segmentation from the previous iteration. The consecutive matching score tells us how close the results are for consecutive iterations. This information can be used in the multi-iteration

process in order to decide when to stop. A complete set of plots of the segmentation scores and the consecutive matching scores for the first 50 iterations, for each individual of the 10 test images, is included in Appendix A.

5.6 Discussion

It is worth pointing out that for each of the 10 images, MMS algorithm performs better than the current OSA for every single instance. We observe the most improvement in iteration 1. We also observe that the segmentation score is even better in few subsequent iterations of MMS. The fact that the results are better in every case, makes the results statistically significant.

The histograms corresponding to the first and sixth iteration of MMS, shown in Figure 5.4(d) and (f), have a bi-modal characteristic. This characteristic shows the separation of bright and dark features. Subsequent iterations of MMS result in a slightly smoother histogram. Some pixels are over-enhanced by MMS, and their grey values are clipped to the maximum of 255. This is indicated by the peaks in the two histograms corresponding to MMS. These over-enhanced pixels often correspond to very bright features in the enhanced objects, and are classified as belonging to oil sand fragments by the subsequent segmentation algorithm.

Image B is handled the best out of the 10 test images by OSA (with segmentation score of 0.6413), and therefore hardly receives any further improvement by MMS. According to Table 5.3, the ground-truth for that image contains the least number of fragments out of the 10 test images, 171, along with Image F. This explains why the results for multiple iterations of MMS fluctuate for Image B and Image F. The score for the initial OSA iteration for Image F, 0.5877, is not as high as for Image B, 0.6413. In addition OSA reports more fragments for Image F, 317, than for Image B, 281. This is evidence that Image F is more difficult to handle by OSA than Image B.

A problem with the multi-iteration approach is not knowing when to stop. One approach to resolve that is to take the consecutive matching score into account. One could monitor the amount of change between consecutive iterations. Once that change becomes small enough, that can be a good place to stop.

5.7 Summary

This chapter presents the experiments and results in order to see how much better the MMS algorithm is in comparison to the existing method used in the OSA. The goal of the contrast enhancement algorithm is to enhance an input image so that it results in improved segmentation, therefore the new algorithm is evaluated in terms of segmentation

performance.

In order to evaluate the performance of MMS, we use a set of 10 manually generated ground-truth images. Each of the ground-truth images is binary, with all the objects in the image delineated from one another. The ground-truth images allow quantitative evaluation of segmentation performance.

The quantitative evaluation of segmentation is computed using a scoring metric, where a score in the range between 0 and 1 is obtained. This score indicates how closely two binary segmented images are matched. The segmented images in our experiments are compared to the manually produced ground-truth images via the score metric.

The experiment is set up so that the initial iteration is performed with the existing OSA algorithm, which produces a priori image required by the first iteration of MMS. Subsequent iterations of the contrast enhancement algorithm are performed, where the segmentation result for a particular iteration is used as a priori data for the next iteration.

The results of the experiment show an average improvement of 0.0518 in the first iteration of MMS in comparison with OSA, which makes that an average relative improvement of 0.0927. An improvement using MMS is observed for each of the ten images, which implies that the results are statistically significant. Image B is enhanced the best by the initial iteration of OSA, and only a minimal improvement is done by MMS. Images B and F have the least number of fragments, according to the ground truth, and the scores for these two images in multiple iterations fluctuate the most.

While the most improvement occurs in the first iteration, a few subsequent iterations produce further improvement. The maximum average segmentation score, within the first 50 iterations, is achieved in the sixth iteration, where the improvement of MMS over OSA is 0.0679, which makes it relative improvement of 0.1223. The average potential captured is 0.2523 for the first iteration, and 0.3271 for the sixth iteration. This potential is measured in comparison to the ground truth provided. Consecutive matching score can be used in a multi-iteration approach in order to decide after which iteration to stop.

Chapter 6

Conclusions

The previous chapter describes the experiments performed to evaluate the MMS algorithm, talks about the evaluation criteria and presents the results of this experiment. In this chapter we provide the conclusion of this thesis. We discuss the lessons learned, and state the objectives that have been achieved. We also provide suggestions for possible improvements in the future.

6.1 Summary

The new contrast enhancement method, MMS, addresses various issues and tries to solve some of the problems present in the existing OSA segmentation method. The multi-scale aspect of this method addresses scale properties of the various objects in the image, where different sizes are dealt with. Another advantage of the scale-space decomposition is the separation of features corresponding to entire object contours, versus internal object texture. Since the texture is detrimental in the segmentation process, it is advantageous to enhance the objects more, without enhancing the texture.

The direct aspect of the enhancement, allows the actual amount of contrast to be controlled. This prevents undesirable artifacts from over-enhancing. The adaptive aspect of this algorithm allows for a separate treatment of individual local regions, thus optimizing the enhancement for each local region.

As seen in the experiments, the new contrast enhancement method results in improved segmentation performance of oil sand images. This is demonstrated by the relative average improvement in segmentation by 0.0927 in the first iteration, and by 0.1223 in the sixth iteration. Based on the presented evidence, we conclude that the new contrast enhancement algorithm improves the performance of the oil sand image segmentation.

The objectives of this thesis are met, where we make a number of research contributions. We develop a new contrast enhancement algorithm for the enhancement of oil sand images. This enhancement algorithm is multi-scale, direct and adaptive. We also show the rela-

tionship between object size and scale of enhancement. Finally, with the empirical results of the experiments, we demonstrate that adequate contrast enhancement improves image segmentation.

6.2 Limitations

There is a number of limitations of this study, both with the MMS algorithm itself, as well as with the way the experiments have been conducted. MMS suffers from a number of limitations. The algorithm is sensitive to provided a priori knowledge. If the initial iteration of OSA makes a significant mistake in segmentation, MMS might not be able to fully recover from it. MMS has a handful of parameters, which require tuning. In addition, the algorithm itself is computationally expansive, where majority of the time is spent in the top-hat multi-scale decomposition.

There is also a number of limitations with the experiments that were conducted. Only ten sample images with corresponding ground truth are used. Having limited number of sample images does not allow for the robustness of the algorithm to be sufficiently tested. In addition the ground truth is biased toward coarser objects, since it does not contain very small objects. This bias has an influence on the results of the experiments.

The ground truth itself is subjective, since it is produced manually. A different person could come up with a different ground truth. The inconsistency is influenced by numerous factors, such as inability to see edges clearly, fatigue and even the characteristics of the display device used in producing the ground truth.

6.3 Future Work

While the MMS algorithm improves the segmentation of oil sand images, there is a number of improvements suggested for future work. The multi-scale decomposition part of this algorithm takes considerable amount of time. In our implementation we decompose the image using equal scale intervals for the entire scale-space.

Due to the fact that the distribution of information in coarser scales has a larger spread over the scales than in finer scales, larger increments can be done in coarser scales and finer in the finest scales. Considering that most of the time is spent in morphological operations using large structuring element, this optimization should introduce a significant reduction in computational time.

Another suggestion for improvement is in the tuning of the parameters. The parameters for this enhancement method are currently tuned on a trial and error basis. A better fine tuning method for optimizing the parameters might be desired, which is more accurate, quicker and automated. This method could be based on concepts from machine learning.

For even more optimal enhancement, a separate set of parameter each individual iteration of MMS might be desired.

Another suggestion of improvement is in the multi-iteration approach of this method. For an optimal and automated functionality, one idea is to make the algorithm decide when to stop, that is after which iteration. One suggestion for finding stopping criteria for the multiple iterations is to monitor the changes in segmentation score between consecutive iterations, for example using the consecutive matching score, and stop when these changes satisfy given criteria.

Bibliography

- [1] B.D. Anderson. A simple technique to determine the size distribution of nuclear crater fallback and ejecta. In *Proc. Symposium on Engineering with Nuclear Explosives*, volume 2, pages 1726–1745, January 14–16, 1970.
- [2] Azeddine Beghdadi and Alain Le Negrate. Contrast enhancement technique based on local detection of edges. *Computer Vision, Graphics, and Image Processing*, 46:162–174, 1989.
- [3] Claude Cunningham. Lessons from the compaphoto technique of fragmentation measurement. In *Proc. FRAGBLAST 5 Workshop on Measurement of Blast Fragmentation*, pages 53–57, August 23–24, 1996.
- [4] F.P. De Vries. Automatic, adaptive, brightness independent contrast enhancement. *Signal Processing*, 21:169–182, October 1990.
- [5] F. Dornaika and H. Zhang. Granulometry with mathematical morphology and motion. In *Proc. 2003 International Conference on Control Science and Engineering*, December 18–20, 2003.
- [6] Wikipedia, The Free Encyclopedia. Contrast, 2006. <http://en.wikipedia.org/wiki/Contrast>.
- [7] John A. Franklin. Fragment size measurements and statistics. In *Proc. FRAGBLAST 5 Workshop on Measurement of Blast Fragmentation*, pages 23–31, August 23–24, 1996.
- [8] Werner Frei. Image enhancement by histogram hyperbolization. *Computer Graphics and Image Processing*, 6:286–294, 1977.
- [9] K.K. Girdner, J.M. Kemeny, A. Srikant, and R. McGill. The split system for analyzing the size distribution of fragmented rock. In *Proc. FRAGBLAST 5 Workshop on Measurement of Blast Fragmentation*, pages 101–108, August 23–24, 1996.
- [10] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Prentice-Hall, Upper Saddle River, NJ, USA, second edition, 2002.
- [11] R. Gordon and R.M. Rangayyan. Feature enhancement of film mammograms using fixed and adaptive neighborhoods. *Applied Optics*, 23(4):560–564, February 1984.
- [12] Ernest L. Hall. Almost uniform distribution for computer image enhancement. *IEEE Transactions on Computers*, C-23(2):207–208, 1974.
- [13] R.A Hummel. Image enhancement by histogram transformation. *Computer Graphics and Image Processing*, 6:184–195, 1977.
- [14] G.C. Hunter, C. McDermott, N.J. Miles, A. Singh, and M.J. Scoble. A review of image analysis techniques for measuring blast fragmentation. *Mining Science and Technology*, 11(1):19–36, 1990.
- [15] Y. Jin, L. Fayad, and A. Laine. Contrast enhancement by multi-scale adaptive histogram equalization. In *Proc. SPIE*, volume 4478, pages 206–213, 2001.
- [16] D.J. Ketcham. Real-time image enhancement techniques. In *Proc. SPIE/OSA*, volume 74, pages 120–125, 1976.
- [17] Syncrude Canada Ltd. website, 2006. <http://www.syncrude.ca/>.

- [18] Norbert H. Maerz, Tom C. Palangio, and John A. Franklin. Wipfrag image based granulometry system. In *Proc. FRAGBLAST 5 Workshop on Measurement of Blast Fragmentation*, pages 91–99, August 23–24, 1996.
- [19] D. Mukherjee and Chatterji B.N. Adaptive neighborhood extended contrast enhancement and its modifications. *Graphical Models and Image Processing*, 57(3):254–265, May 1995.
- [20] S. Mukhopadhyay and B. Chanda. A multiscale morphological approach to local contrast enhancement. *Signal Processing*, 80(4):685–696, April 2000.
- [21] N. Otsu. A threshold selection method from gray level histogram. *IEEE Transactions on Systems, Man and Cybernetics (SMC)*, 9:62–66, 1979.
- [22] R.B. Paranjape, W.M. Morrow, and R.M. Rangayyan. Adaptive-neighborhood histogram equalization for image enhancement. *CVGIP: Graphical Models and Image Processing*, 54(3):259–267, 1992.
- [23] T. Peli and J.S. Lim. Adaptive filtering for image enhancement. *Optical Engineering*, 21(1):108–112, January–February 1982.
- [24] William B. Pennebaker and Joan L. Mitchell. *JPEG Still Image Data Compression Standard*. Kluwer Academic Publishers, Norwell, MA, USA, 1992.
- [25] S.M. Pizer, E.P. Amburn, J.D. Austin, R. Cromartie, A. Geselowitz, T. Geer, B.H. Romeny, J.B. Zimmerman, and K. Zuiderveld. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3):355–368, September 1987.
- [26] Stephen M. Pizer. Intensity mappings for the display of medical images. In *Proc. 11th Annual Symposium on The Sharing of Computer Programs and Technology in Nuclear Medicine*, pages 205–217, 1981.
- [27] Mark Polak, Hong Zhang, and Minghong Pi. Objective evaluation metric for image segmentation of multiple objects for size measurement. Technical report, University of Alberta, 2006.
- [28] Andrea Polesel, Giovanni Ramponi, and V. John Mathews. Image enhancement via adaptive unsharp masking. *IEEE Transactions on Image Processing*, 9(3):505–510, 2000.
- [29] Kelly Rehm and William J. Dallas. Artifact suppression in digital chest radiographs. In *Proc. SPIE*, volume 1092, pages 290–300, 1989.
- [30] J.C. Santamarina, M. Morley, J.A. Franklin, and D.S. Wang. Development and testing of a zooming technique for fragmentation measurement. In *Proc. FRAGBLAST 5 Workshop on Measurement of Blast Fragmentation*, pages 133–139, August 23–24, 1996.
- [31] J. Tang, E. Peli, and S. Acton. Image enhancement using a contrast measure in the compressed domain. *IEEE Signal Processing Letters*, 10(10):289–292, October 2003.
- [32] A. Toet. Adaptive multi-scale contrast enhancement through non-linear pyramid recombination. *Pattern Recognition Letters*, 11(11):735–742, November 1990.
- [33] Marc Van Droogenbroeck and Hugues Talbot. Fast computation of morphological operations with arbitrary structuring elements. *Pattern Recognition Letters*, 17:1451–1460, 1996.
- [34] Bernard Widrow and Samuel D. Stearns. *Adaptive Signal Processing*. Prentice-Hall, Englewood Cliffs, NJ, USA, 1985.
- [35] Andrzej Zadorożny, Hong Zhang, and Martin Jägersand. Granulometry using image transformation techniques. In *Proc. Vision Interface*, pages 433–438, May 2002.

Appendix A

Individual Experiment Results

A.1 Result Images and Plots

This appendix presents individual results for each of the ten images used in the experiments. Figures A.1-A.10 show the resultant images and their corresponding segmentations. These images are obtained for the initial OSA run, the first iteration of MMS and the iteration of MMS with the maximum segmentation score. The input images and the corresponding ground truth images are also included.

Figures A.11-A.20 show fragment size distributions in terms of cumulative percentage passing plots. These percentage passing plots are calculated by area. These plots are obtained for the initial run of OSA, the first iteration of MSS, and the iteration of MSS with highest segmentation score. In all these plots the cumulative percentage of corresponding ground-truth images is included as well.

Figures A.21-A.30 show plots of the segmentation scores and the consecutive matching scores, using the first 50 iterations of the MMS algorithm. Segmentation scores are computed by comparing the output segmented image from a given iteration of the algorithm to the corresponding ground-truth image. Consecutive matching scores are computed by comparing the output segmented image from a given iteration to the output segmented image from the previous iteration. Iteration 0 refers to the output of OSA, which provides the initial a priori data. The subsequent iterations are of the MMS algorithm.

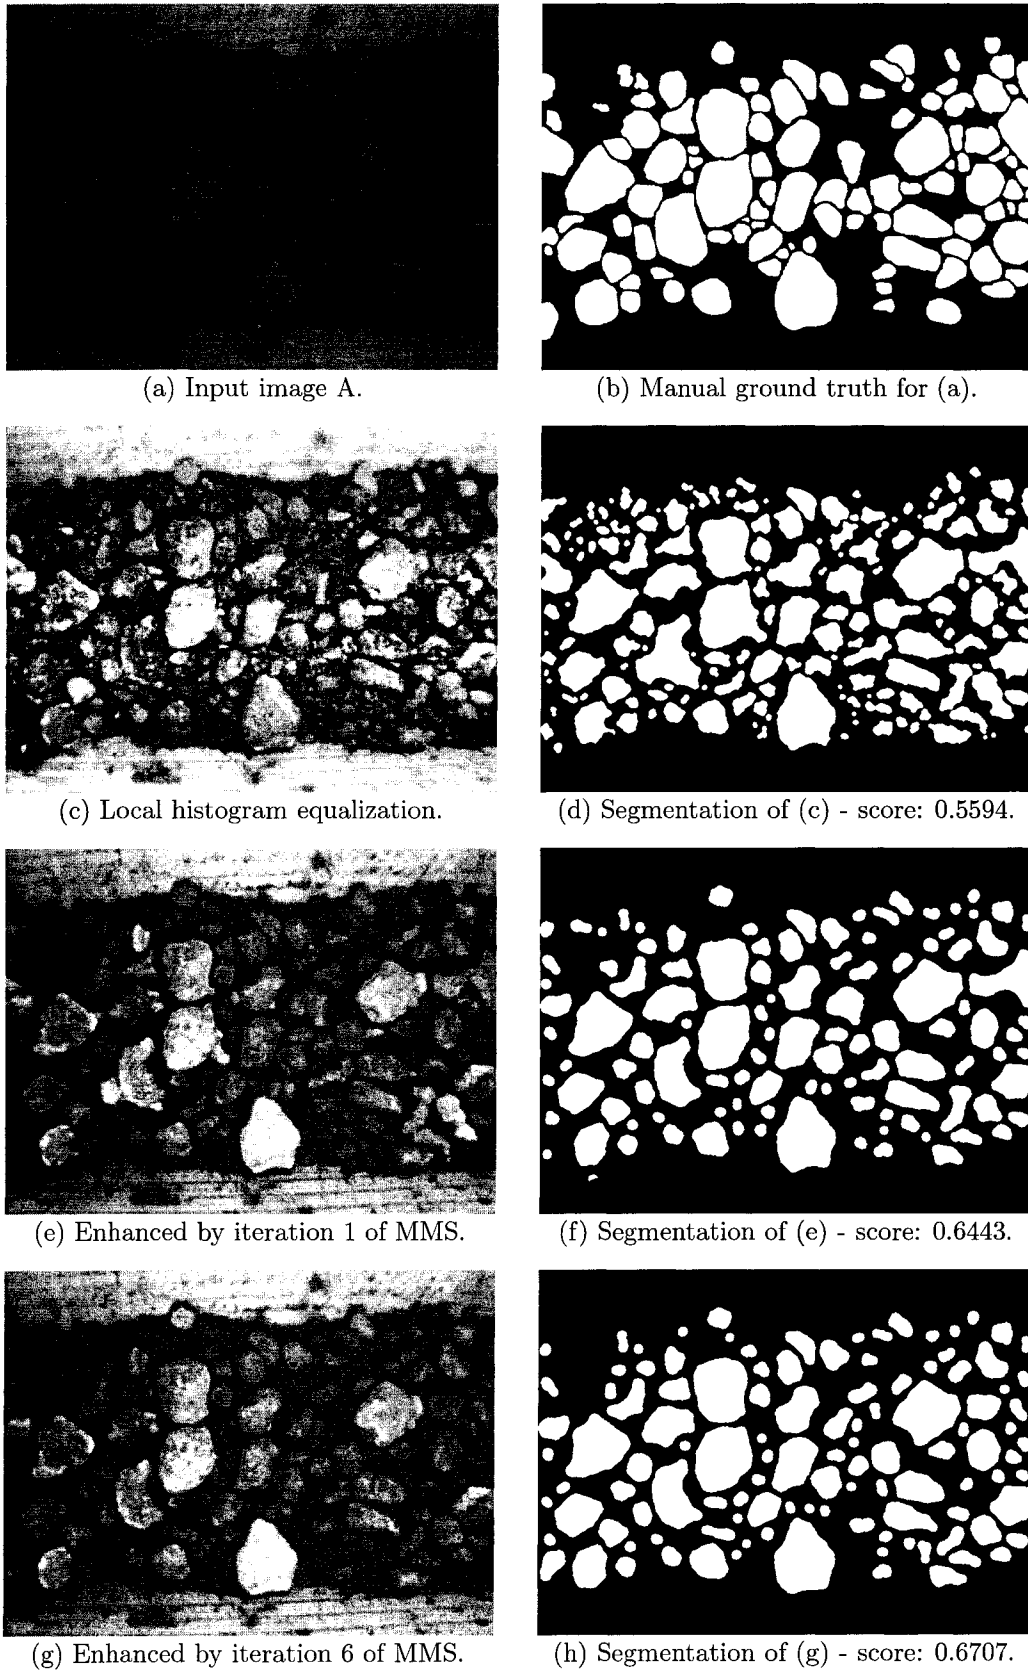



Figure A.1: Enhancement and segmentation results for Image A.




(a) Input image B.



(b) Manual ground truth for (a).




(c) Local histogram equalization.



(d) Segmentation of (c) - score: 0.6413.




(e) Enhanced by iteration 1 of MMS.



(f) Segmentation of (e) - score: 0.6487.




(g) Enhanced by iteration 25 of MMS.




(h) Segmentation of (g) - score: 0.6737.

Figure A.2: Enhancement and segmentation results for Image B.



(a) Input image C.



(b) Manual ground truth for (a).




(c) Local histogram equalization.



(d) Segmentation of (c) - score: 0.5873.




(e) Enhanced by iteration 1 of MMS.



(f) Segmentation of (e) - score: 0.6476.




(g) Enhanced by iteration 3 of MMS.




(h) Segmentation of (g) - score: 0.6667.

Figure A.3: Enhancement and segmentation results for Image C.




(a) Input image D.



(b) Manual ground truth for (a).




(c) Local histogram equalization.



(d) Segmentation of (c) - score: 0.4952.




(e) Enhanced by iteration 1 of MMS.



(f) Segmentation of (e) - score: 0.5341.




(g) Enhanced by iteration 17 of MMS.




(h) Segmentation of (g) - score: 0.5651.

Figure A.4: Enhancement and segmentation results for Image D.




(a) Input image E.




(b) Manual ground truth for (a).




(c) Local histogram equalization.




(d) Segmentation of (c) - score: 0.5909.




(e) Enhanced by iteration 1 of MMS.



(f) Segmentation of (e) - score: 0.6347.




(g) Enhanced by iteration 4 of MMS.




(h) Segmentation of (g) - score: 0.6565.

Figure A.5: Enhancement and segmentation results for Image E.




(a) Input image F.




(b) Manual ground truth for (a).




(c) Local histogram equalization.




(d) Segmentation of (c) - score: 0.5877.




(e) Enhanced by iteration 1 of MMS.



(f) Segmentation of (e) - score: 0.6259.



(g) Enhanced by iteration 3 of MMS.



(h) Segmentation of (g) - score: 0.6553.

Figure A.6: Enhancement and segmentation results for Image F.

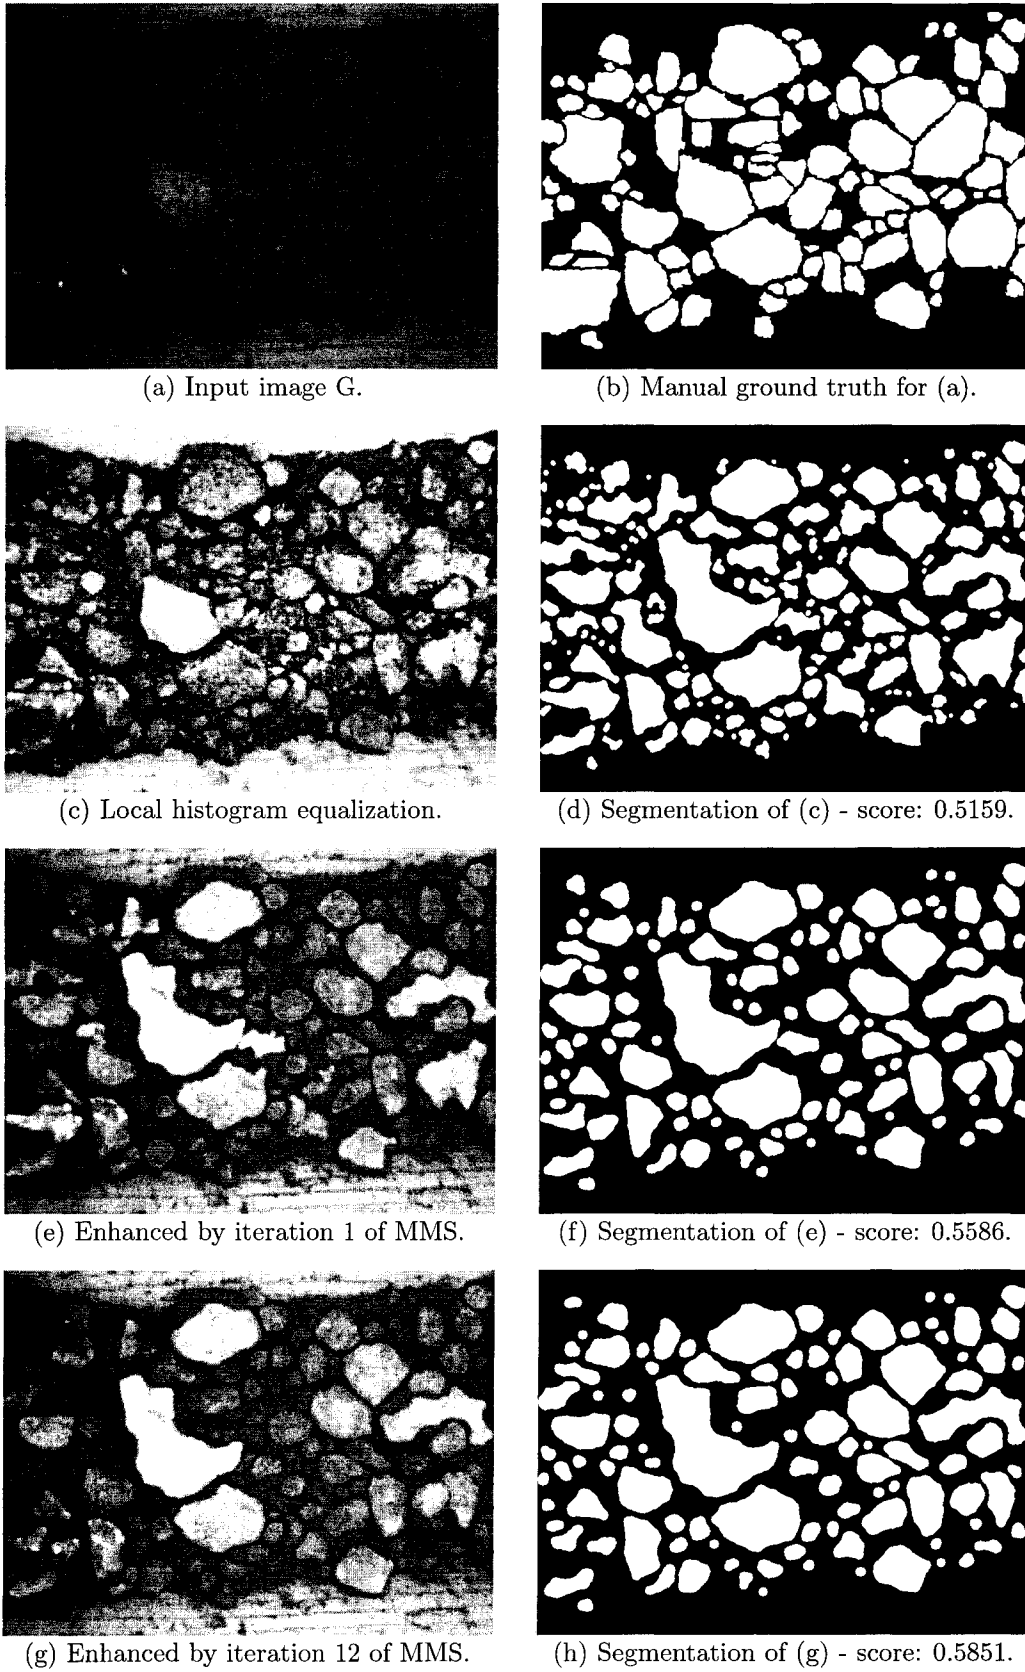




Figure A.7: Enhancement and segmentation results for Image G.



(a) Input image H.



(b) Manual ground truth for (a).




(c) Local histogram equalization.




(d) Segmentation of (c) - score: 0.5810.




(e) Enhanced by iteration 1 of MMS.



(f) Segmentation of (e) - score: 0.6467.



(g) Enhanced by iteration 11 of MMS.



(h) Segmentation of (g) - score: 0.6690.

Figure A.8: Enhancement and segmentation results for Image H.







Figure A.9: Enhancement and segmentation results for Image I.



(a) Input image J.




(b) Manual ground truth for (a).




(c) Local histogram equalization.




(d) Segmentation of (c) - score: 0.5543.




(e) Enhanced by iteration 1 of MMS.



(f) Segmentation of (e) - score: 0.6223.




(g) Enhanced by iteration 30 of MMS.




(h) Segmentation of (g) - score: 0.6442.

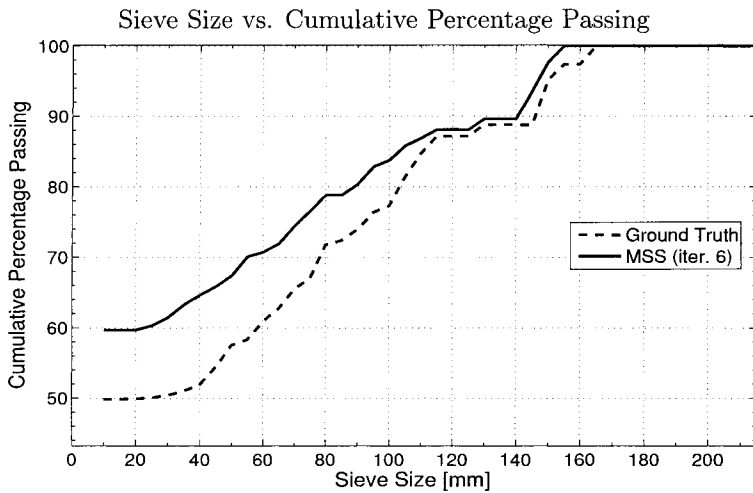
Figure A.10: Enhancement and segmentation results for Image J.



(a) Cumulative percentage passing for the initial iteration of OSA (329 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (219 fragments).




(c) Cumulative percentage passing for iteration 6 of MSS (227 fragments).


Figure A.11: Cumulative percentage passing for Image A (213 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (281 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (187 fragments).




(c) Cumulative percentage passing for iteration 25 of MSS (187 fragments).


Figure A.12: Cumulative percentage passing for Image B (171 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (347 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (233 fragments).




(c) Cumulative percentage passing for iteration 3 of MSS (239 fragments).


Figure A.13: Cumulative percentage passing for Image C (237 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (307 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (225 fragments).




(c) Cumulative percentage passing for iteration 17 of MSS (237 fragments).


Figure A.14: Cumulative percentage passing for Image D (213 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (347 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (275 fragments).




(c) Cumulative percentage passing for iteration 4 of MSS (285 fragments).


Figure A.15: Cumulative percentage passing for Image E (303 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (317 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (211 fragments).




(c) Cumulative percentage passing for iteration 3 of MSS (203 fragments).


Figure A.16: Cumulative percentage passing for Image F (171 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (315 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (215 fragments).




(c) Cumulative percentage passing for iteration 12 of MSS (205 fragments).


Figure A.17: Cumulative percentage passing for Image G (207 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (337 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (199 fragments).




(c) Cumulative percentage passing for iteration 11 of MSS (211 fragments).


Figure A.18: Cumulative percentage passing for Image H (213 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (363 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (273 fragments).




(c) Cumulative percentage passing for iteration 6 of MSS (261 fragments).


Figure A.19: Cumulative percentage passing for Image I (273 fragments in ground truth).



(a) Cumulative percentage passing for the initial iteration of OSA (375 fragments).




(b) Cumulative percentage passing for the first iteration of MSS (289 fragments).




(c) Cumulative percentage passing for iteration 30 of MSS (289 fragments).

Figure A.20: Cumulative percentage passing for Image J (309 fragments in ground truth).




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.21: Results of the first 50 iterations of MMS on Image A.




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.22: Results of the first 50 iterations of MMS on Image B.




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.23: Results of the first 50 iterations of MMS on Image C.




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.24: Results of the first 50 iterations of MMS on Image D.




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.25: Results of the first 50 iterations of MMS on Image E.




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.26: Results of the first 50 iterations of MMS on Image F.




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.27: Results of the first 50 iterations of MMS on Image G.




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.28: Results of the first 50 iterations of MMS on Image H.




(a) Segmentation score for the first 50 iterations of MMS.




(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.29: Results of the first 50 iterations of MMS on Image I.



(a) Segmentation score for the first 50 iterations of MMS.



(b) Consecutive matching score for the first 50 iterations of MMS.

Figure A.30: Results of the first 50 iterations of MMS on Image J.