SEMI-AUTOMATED CHARACTERIZATION OF THIN-SECTION PETROGRAPHIC IMAGES

DARRELL MOULAND

Semi-Automated Characterization of Thin-Section Petrographic Images

By

©Darrell Mouland, B.Eng.

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfillment
of the Requirements of the Degree of
Master of Engineering

Faculty of Engineering and Applied Science Memorial University of Newfoundland

May 2005

St. John's Newfoundland

Abstract

This thesis represents the continuation of work on PetrograFX, an automated image analysis toolkit for petrographic image analysis. These types of images are used in the petroleum industry to provide valuable information, however, the retrieval of data from these images is time consuming and prone to operator bias. An integrated solution that combines a number of basic image processing concepts, each tailored towards segmenting a particular type of grain, is developed to automate this process. Specifically, an attempt is made to replicate the methodology and analysis carried out by core laboratories, which typically place more emphasis on overall interpretation of the image rather than just the measurement of the porosity and quartz grain distribution. This requires a solid treatment of the geological background to ensure the data being collected will be useful. Due to their complex nature there will be regions within these images that are unidentifiable. This approach necessitates a classification routine to eliminate objects once they have been segmented to ensure that they are unaffected by subsequent routines. To provide a quick and objective assessment segmentation performance an automated accuracy routine is presented.

Acknowledgements

I would like to take this opportunity thank my supervisors, Dr. Raymond Gosine and Dr. George Mann for providing me with financial assistance and excellent facilities to conduct my research. I would also like to thank my friends and colleagues for their support, most notably Andrew and Angela, for their guidance and ability to keep things in perspective. Finally, I would like to thank Melanie for being such a great person and my family for making this all possible.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

The petroleum industry is multi-disciplined and requires knowledge of both engineering and earth sciences. In recent years technology has had a tremendous affect on production efficiencies. With the advent of reservoir simulators and other computer programs the industry is moving away from the qualitative constraints of traditional geological and petrophysical practices and is becoming evermore quantified and precise. Despite these advances the fact remains that the core, collected during the drilling of a well, remains the only direct source of data for the reservoir. The geologist uses information acquired from the core in combination with other sources such as well logs and seismic data to assess the potential of a particular reservoir. Coring a well can be an expensive procedure and is, therefore, not carried out for every well. This greatly increases the importance of the data that are obtained, especially considering the fact that wells can be several kilometers apart and the well data must be extrapolated over this distance. Core analysis involves testing samples of the core for rock properties such as porosity and permeability. The benefit of core analysis is that it is unambiguous and provides an accurate interpretation of the rock structure, however, due to the limited sampling along the length of the well, it cannot properly detect heterogeneities. Well logs, generated by passing various sensors and instruments through the well bore in the area of the formation, have the advantage of being able to scan the entire length of the well with an effective radius ranging from a few inches to a couple of feet. These logs are, therefore,

better suited to detect reservoir heterogeneities but they require assumptions about properties of the rock structure as well as about any fluids that may be present in the reservoir [Koederitz et al., 1989]. While each method provides satisfactory results under favorable conditions (homogeneous reservoir), integration of all available data into one cohesive model is the key to successful reservoir analysis in real world situations.

1.1 Petrographic Image Analysis

Petrographic image analysis (PIA) is one type of core analysis in which information is acquired by producing thin sections of the core that are then viewed under a microscope. These thin sections provide a good qualitative indication of the reservoir properties by allowing the geologists to visualize the grains, pores and other materials. Thin-section acquisition is not limited to conventional core samples; they can also be acquired from sidewall cores and cuttings samples, which can be important in cases where conventional core collection is limited [Davies, 1990]. Pore space data acquired from thin-sections can be used to estimate permeability and predict reservoir quality based on an understanding of diagenetic alteration and its relation to effective porosity [Ruzyla, 1986; Bowers et al., 1994; Gies and McGovern, 1993; James, 1995; Mowers and Budd, 1996]. In addition, this information can also be linked to well log data [Davies, 1990].

It is difficult to obtain accurate quantitative measures from PIA since a highly skilled operator is required to analyze a large number of grains (300-500) to make the results statistically meaningful [Kennedy and Mazzullo, 1991]. This type of analysis can be very time consuming and it can take 2-3 hours to complete a slide containing 200-500 grains [Gosine and Burden, 1999], making it uneconomical to complete more than a few

images per well. In addition to this, a standard deviation of 5% between operators can be expected when manually segmenting images [Diogenes et al., 2003]. If a greater number of these images could be analyzed with greater accuracy, more information would be available to integrate into the reservoir model. To facilitate this, efforts have been extended to characterize thin sections using automated image analysis techniques. Fully automated image analysis has been complicated by the presence of touching, overlapping, and fused grains that lack distinct boundaries. An inability to separate these grains using image-processing methods leads to substantial bias when calculating textural characteristics of the sample [van de Berg et al., 2002]. Several image-processing algorithms have been proposed to remedy this situation but when it comes to accurately segmenting the entire image, results are generally poor to moderate. Difficulties with these methods have prompted researchers to focus their attention on developing alternative methods for acquiring the thin-section images. For some image acquisition methods, the visibility of grain boundaries vary depending on the orientation of the sample, therefore, a composite image can be formed which contains significantly more boundaries as compared to other methods [Starkey and Samantaray, 1994; Ross et al., 2001]. Despite the above problems, most thin-section images are still acquired using more traditional methods and given the fact that many oil companies may already have an inventory of these images, it is worthwhile to pursue automated solutions based on these images. In addition, these methods can be adapted for use with other types of samples and materials, including metallurgical samples obtained using reflected light sources [Kraatz et al., 2003].

1.2 Motivation and Objectives

The research presented here is the continuation of PetrograFX, an automated thin section image analysis software package developed by C-Core, Hibernia Management and Development Company (HMDC), the Center for Earth Resources Research and Memorial University of Newfoundland. The goal of this software is to produce pore space and quartz grain segmentation and measurement algorithms that are capable of automated, fast and accurate thin-section image analysis to either replace or partially replace the human analysis [Zhao, 2000]. Work completed prior to the start of this thesis included an automatic pore space segmentation and measurement algorithm and a quartz grain segmentation and measurement algorithm. The algorithm for measuring the pore space was well developed and provided excellent results, while the routine for quartz grain segmentation simply provided a starting point for further research. It is recognized that this routine works well for segmenting a certain of grain, however, there are many other grains on which it fails.

The overall goal of this thesis is to provide additional methods for segmenting the remaining grains, as well as identifying additional image constituents. Specifically, an attempt is made to replicate the methodology and analysis carried out by core laboratories, which typically place more emphasis on overall interpretation of the image rather than just the measurement of the porosity and quartz grain distribution. This requires a solid treatment of the geological background to ensure that the data being collected will be useful. It is important to recognize that because constant operator interaction is required to properly identify all constituents of thin-section images [Starkey and Samantaray, 1994], there will be regions within these fixed images that are

unidentifiable. In addition, the difference in interpretation that exists among experienced operators suggests that, even with constant interaction with the sample, some regions are impossible to segment accurately. Russ [1990] points out that this subjectivity is unavoidable when acquiring measurements using microscopy. Based upon this information, distinct types of grains were grouped together and different methods were developed to segment different grain types, with some types remaining un-segmented due to lack of information or uncertainty in the interpretation. The theory is that enough of the image can be segmented so as to provide accurate information for image classification (in a geological sense). This differs from previous methods that applied one segmentation routine to the entire image [Russ, 1995; Zhao, 2000; van den Berg, 2002], but is similar to genetic algorithms that evolve a different routine for each mineral in the sample [Ross et al., 2001]. This approach necessitates a classification routine to eliminate objects once they have been segmented to ensure that they are unaffected by subsequent routines.

Another important factor to consider is that the grains seen in these images are the result of natural processes and can therefore show significant variation from image to image. This variation further complicates the design of an image analysis algorithm; therefore, an approach that includes more than one segmentation routine provides adaptability from image to image.

To gauge the effectiveness of the segmentation algorithm, a quick and objective method is needed, otherwise quantifying the performance will suffer from the same human subjectivity that was presented earlier. To accomplish this, potential grains are to be compared automatically to their manually segmented counterparts.

To summarize, the primary goal is to develop an adaptable segmentation algorithm capable of accurately identifying a sufficient amount of material to provide geological classification for a range of images. This is accomplished by developing a modular design that integrates a number of approaches, each aimed at a specific type of grain material. Secondary goals include identification of other image constituents and the development of automated methods to quantify the performance of the algorithm. These automated methods, in combination with the modular design of the segmentation algorithm, introduce the potential for learning as the program can test for the parameters and routines that maximize the accuracy of the segmentation.

Chapter 2 provides an overview of thin-section petrographic image analysis and the research that has been carried out in this area. Chapter 3 discusses the manual analysis of the test images and the specific features found in the data set. Chapter 4 outlines the development of the image processing methods used to segment, classify and measure the image constituents. Chapter 5 quantifies the performance of the segmentation routines and discusses their integration to form a single segmentation algorithm. Chapter 6 presents and discusses the results of the integrated segmentation routine. Chapter 7 presents the conclusions and recommendations.

Chapter 2

Petrographic Image Analysis

For any image-processing application, it is necessary to have a basic understanding of the contents of the images under consideration. In other words, it is necessary to know what is being observed and why. In the case of petrographic images, the contents are extremely varied due to their complex origins. Typically, a number of different image components are examined to provide some higher-level geological 'understanding' or interpretation. The goal of this chapter is to introduce the geological concepts required to identify and measure the contents of the petrographic images and to provide a review of previous work completed in this field.

2.1 Geological Background

Petrology is the branch of geology that is concerned with the study of rocks, including igneous, sedimentary, and metamorphic. Oil reservoirs most commonly consist of sandstone, although other types do exist. The sandstone, containing the oil, is formed through a depositional process that consists of the transportation and early burial of detrital material and allochems followed by the process of diagenesis, which turns the sand into sandstone [Raymond, 2002]. The detrital material consists of pre-existing rock fragments or grains originating from igneous, metamorphic and sedimentary sources. Allochems are any chemical or biochemical material that was formed prior to deposition, i.e., in an earlier rock formation. The term authigenic refers to any part of the rock that

was formed *in situ,* as part of the diagenetic process, including chemical and biochemical precipitates. Generally there are voids within the sandstone known as porosity. Precipitates form in the interstitial pore space, thereby, resulting in a reduction of the primary porosity. Aspects of Sedimentary Petrology are drawn upon to classify these petrographic thin-section images. This classification includes an examination of the texture, as well as the composition and alteration methods, since these factors provide the most important information for determining the origin, method of formation, and overall quality of the reservoir structure.

2.1.1 Sedimentary Composition

The composition refers to the mineralolgy or chemical characteristics of the grains. Compositional classification schemes are based on the relative amounts of quartz, feldspars and rock fragments, since they are the most common minerals forming detrital grains [Lewis and McConchie, 1994b]. These schemes can be very detailed, depending on the material present in the samples and the required depth of analysis. Sandstones are technically composed of grains in the range of 2.0-0.06 mm in diameter, however, textural studies incorporate grains in the range of 4.0-0.03 mm in diameter and so this convention is sometimes adopted for compositional classification. Material smaller than 0.03mm in diameter is classified as silt and clay material and is considered to be part of the matrix (as opposed to part of the grain fabric). Matrix material and authigenetic material (cements) are also considered when identifying the composition of sandstone.

Quartz varieties include volcanic, vein, recycled sedimentary, metamorphic, and common grains [Lewis and McConchie, 1994b]. The volcanic grains are typically water

clear, unstrained and mono-crystalline. Grains of the vein variety tend to have abundant vacuoles or other inclusions and may be unstrained or slightly strained. They can be mono-crystalline or polycrystalline with obscure crystal boundaries. The recycled sedimentary grains tend to show the characteristics of the ultimate source rather than the intermediate sedimentary source, but the presence of rounded or worn overgrowths is the main means of identification. Metamorphic grains are difficult to identify and are often grouped in the common grain category. Common grains are those that do not have any clear indicators of their origin. These grains can be mono-crystalline or poly-crystalline, strained or unstrained, with or without vacuoles or other inclusions.

Feldspar grains are best identified by staining and when present, the discrimination of the different varieties can provide information about provenance. It is difficult to locate these grains without the stains; however, the use of crossed-polarized light may prove useful. Also, other characteristics, such as twinning or the alteration of feldspars to other materials, may be indicators of the presence of a particular variety. In some cases, feldspar grains can be completely replaced by authigenic clays during diagenesis.

Rock fragments are the best indicators of provenance and should be carefully sought in thin section [Lewis and McConchie, 1994b]. However, these are not easily identifiable and in plane-polarized light they may be colorless or a murky pale brown and they are often confused with altered feldspars. Once found, techniques of optical microscopy are able to give further insight into the classification of the different rock fragments but there are no set rules separating the different classes.

Matrix material is composed of clay and silt-sized particles deposited with the other detrital grains and it has the potential to partially or completely fill the interstitial pore spaces. Cements, on the other hand, are precipitated after deposition during the diagenetic process in the form of optically-continuous, and often euhedral, overgrowths on the original detrital grains. Analysis of these materials provides important information on the diagenetic history of the formation.

2.1.2 Sedimentary Texture

In terms of sedimentology, texture refers to sediment fabric (including grain size and sorting) as well as the individual grain shapes. Sediment fabric refers to the packing, Figures 2.1A and B [Tucker, 2001], and orientation of the grains with specific emphasis on the types of contacts experienced between the grains. This is a very important factor for thin-section analysis and a thorough understanding of this factor is required for proper interpretation, and hence segmentation, of grain boundaries. Three main types of contacts exist, namely, sutured, point and concavo-convex, shown in Figure 2.1C, D, and E respectively. In terms of orientation, the grains can either have a preferred orientation, Figure 2.1F, or show no orientation at all. Finally, the fabric can be either grain supported or matrix supported as shown in Figure 2.1G and H. Porosity and permeability are highly dependent upon the packing of the grains and the packing is, in turn, dependent on the grain size and sorting. High porosities generally imply loose packing, approaching that of cubic packing, Figure 2.1A. Lower porosities are the result of tighter packing in the form of the rhombohedral arrangement, Figure 2.1B.

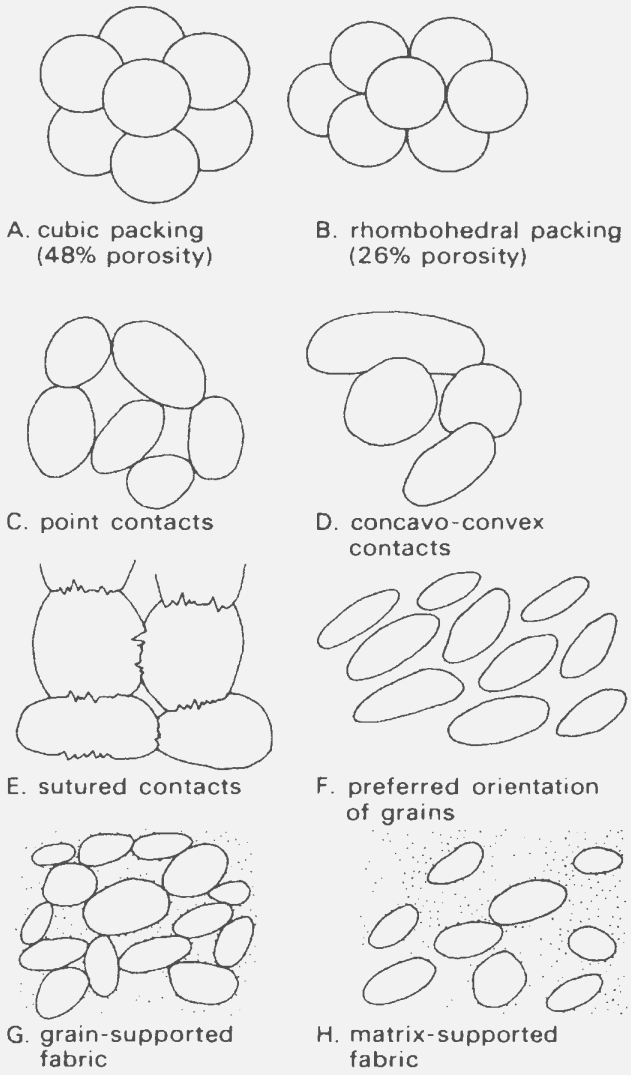

Figure 2.1: Sediment fabric descriptions [Tucker, 2001].

Grain size is typically defined using the Wentworth scale that was originally developed by Udden. This scale is geometric in nature and is divided such that each boundary is either one half or twice the millimeter value of the next boundary. The scale was further modified to include phi values so each boundary is an integer rather than a fraction. This scale system is summarized and shown in Figure 2.2 [Lewis and McConchie, 1994b].

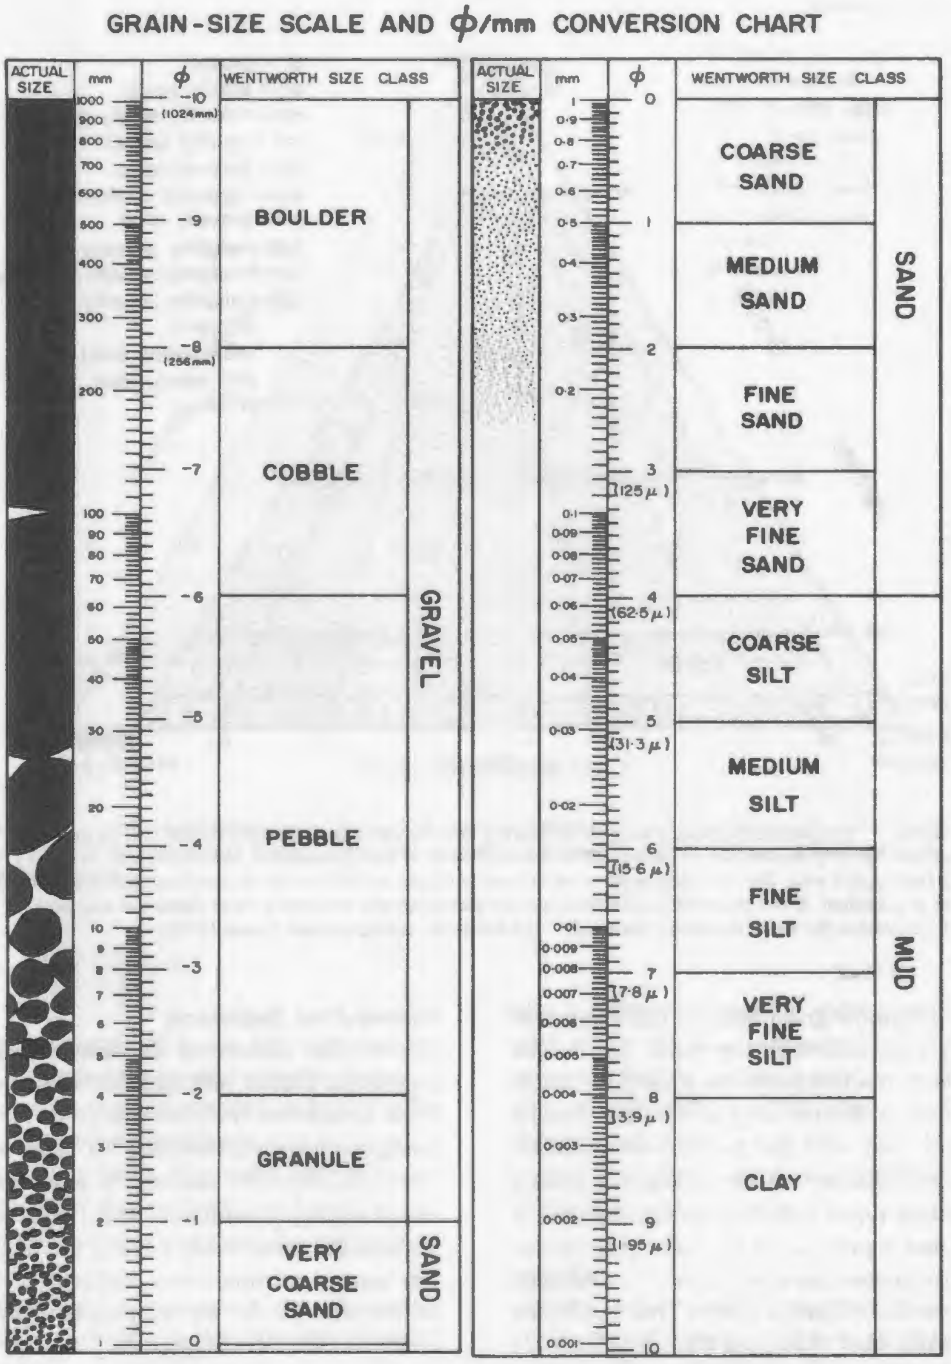

Figure 2.2: Wentworth scale showing grain size classification [Lewis and McConchie, 1994b].

Sorting is a measure of the grain size distribution; poorly sorted samples have a wide range of grain sizes while those that are well sorted have similar grain sizes. Sorting is an

important indicator of depositional history and is also closely related to porosity. Comparative figures for determining sorting are shown in Figure 2.3 [Scholle, 1979].

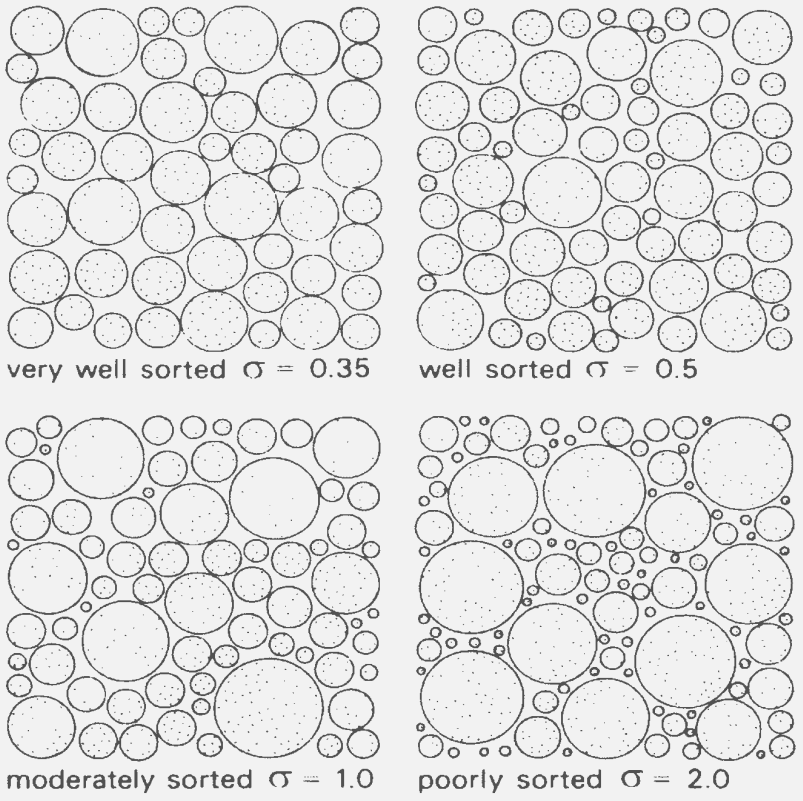

Figure 2.3: Sorting images [Scholle, 1979].

Grain Shape

Most terms used to describe grain shape are related to three-dimensional (3-D) grains as opposed to the two-dimensional (2-D) forms that are seen in thin-section. Two descriptions that have 2-D counterparts are sphericity and surface texture [Raymond, 2002]. Spherecity is a difficult property to measure since the true value is taken to be the ratio of the grain surface area to the surface area of a sphere of the same volume. However, a grain that has high sphericity in 3-D will posses a high degree of circularity in 2-D. Surface texture is not easily observed or measured in thin-section since it is a property of the three-dimensional grain structure but there is some overlap with the

concepts of roundness and angularity. A grain that has a rough or angular surface can be observed in thin-section to be angular as well. For most studies, it is sufficient to use comparative silhouette diagrams for qualitative measures of roundness, as shown in Figure 2.4 [Lewis and McConchie, 1994b]. Six degrees of roundness/angularity are defined.

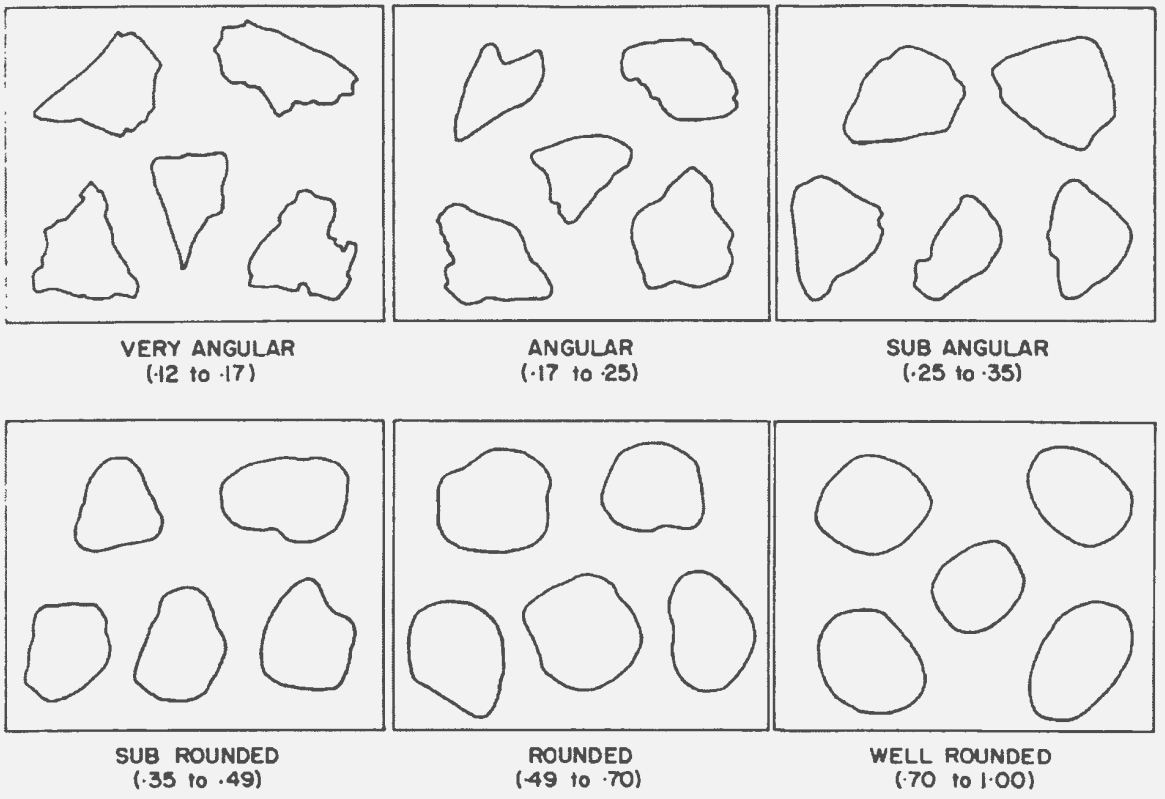

Figure 2.4: Silhouettes used for estimating roundness [Lewis and McConchie, 1994b]

2.1.3 Alteration Methods

Alteration methods are better known as diagenesis, the term used to describe all processes that affect sediments after their final deposition. These processes are very important and can strongly influence the characteristics and the storage potential of a given reservoir. These effects may be negligible or extensive, local or widespread, essentially chemical or mostly physical. They may be obvious in thin-section or they may only be apparent after

analysis with special equipment such as cathodoluminescence, scanning electron microscopy, x-ray diffraction, x-ray fluorescence, or electron probe analysis [Lewis and McConchie, 1994b].

Chemical alteration methods do not generally destroy the dominant detrital components. However, grains of chemically unstable minerals may be completely dissolved or altered and replaced to such an extent that their original character is obscured. For example, feldspars and volcanic rock fragments may be altered to clay materials to such a degree that they become difficult to distinguish from clayey sedimentary rock fragments or clayey matrix. New minerals may be formed during diagenesis by direct precipitation from interstitial fluids. This process of precipitation tends to form euhedral crystals, either attached to the surface of existing grains, or on their own. Factors affecting their growth include the presence of other detrital components that may interfere and prevent the crystals from forming perfect crystal facies and therefore cause a sutured boundary as shown in Figure 2.5 [Sippel, 1968].

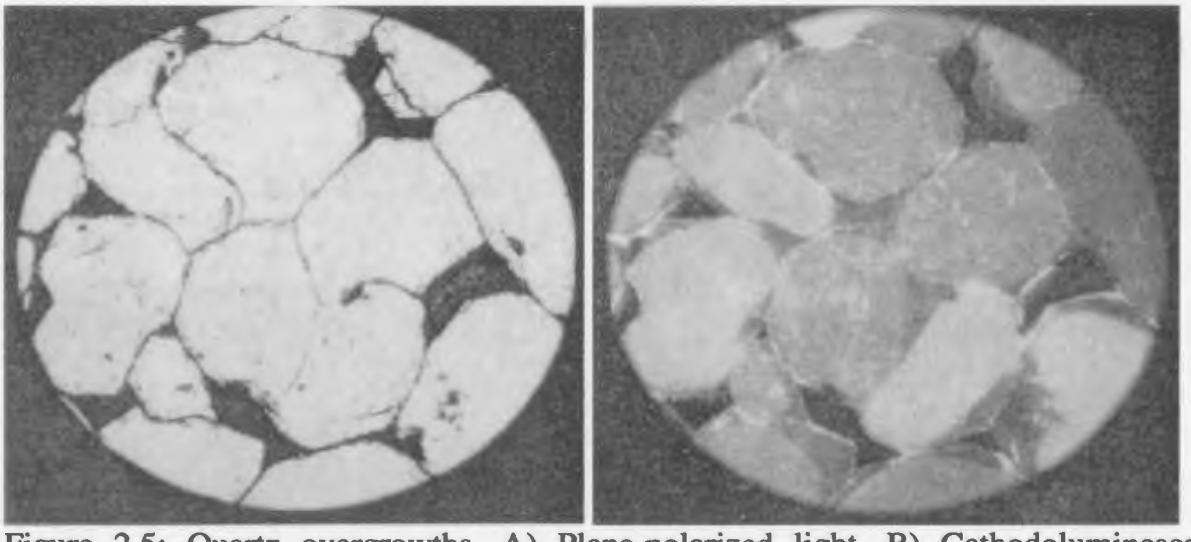

Figure 2.5: Quartz overgrowths. A) Plane-polarized light. B) Cathodoluminescence [Sippel, 1968].

This authigenic material, may either partially or totally, fill the primary interstitial pore space. Physical alteration, in the form of pressure solution, can occur when two grains are touching and pressure causes one or both grains to dissolve preferentially at the points of greatest pressure. This causes the grains to interpenetrate and interlock and can be confused with overgrowth interlocking as shown in Figure 2.6 [Sippel, 1968]. Other physical effects of diagenesis include fracturing of grains and deformation of matrix material.

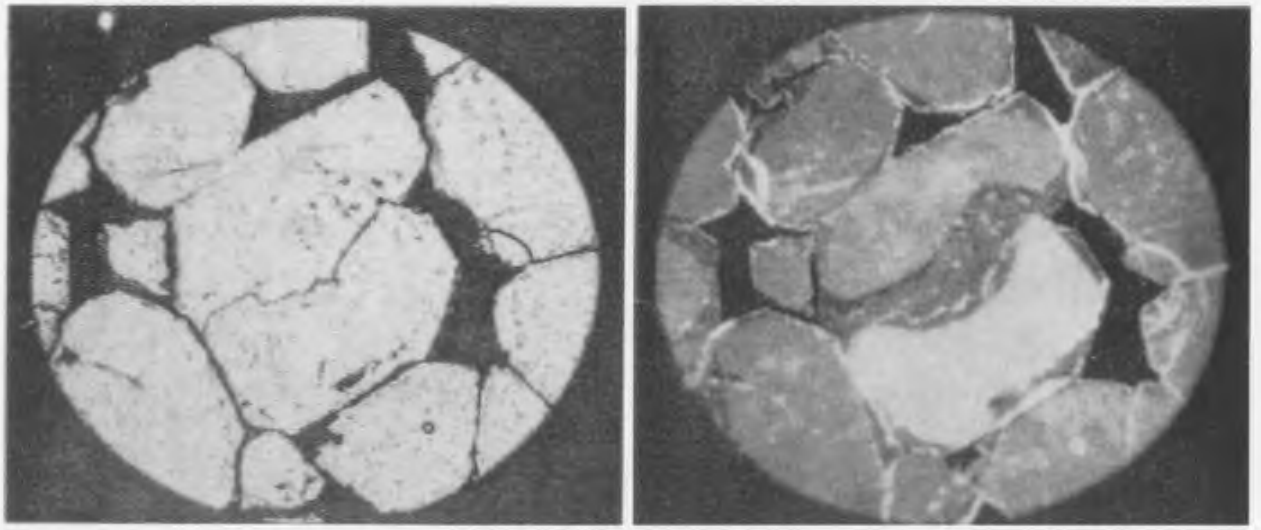

Figure2.6: A) In plane-polarized light the sutured boundary taken as a sign of pressure solution. B) Cathodoluminescence shows it to be the result of interfering quartz overgrowths [Sippel, 1968].

Diagenetic effects can be responsible for substantial modification to grain roundness and grain size. Also, the resulting changes in porosity and permeability caused by these modifications are very important in the migration and entrapment of petroleum resources. In some cases, the primary porosity may be completely destroyed but there is also the potential for secondary porosity to be created.

2.1.4 Acquiring Data

Modal Analysis

Modal Analysis is used to determine the relative amounts of image constituents for compositional classification schemes. This is accomplished by traversing a given slide in a grid type pattern until 300-500 points on mineral grains are acquired [Lewis and McConchie, 1994a]. When selecting acquisition points, particular attention should be paid to samples with heterogeneities. The number frequency obtained using this method is, for most purposes, an adequate estimate of the relative abundance of components. It is important to note that this method suffers from operator bias when identifying the minerals, rock fragments and matrix.

Grain Size and Shape

Grain size can be estimated by using either the equivalent diameter [Kennedy and Mazzullo, 1991] or the length of the long axis of the grains [El-Dein et al., 1984]. Circularity can be estimated by the aspect ratio, defined as the ratio the grains apparent short axis to its apparent long axis [van den Berg et al., 2002]. For studies that require quantitative measures of roundness, the Waddell formula can be used [Tucker, 2001]:

Sorting Measures

Two major methods have been proposed to estimate sorting [Lewis and McConchie, 1994a]. The first is based on graphical methods while the other, known as the method of moments, is computational in nature. Several factors make the method of moments favorable for the determination of sorting parameters. Firstly, it can be computerized

and, therefore, easily implemented into any image processing software. Secondly, it supposedly has advantages over other graphical methods in that it is not biased by the assumption of a fundamental lognormal distribution. This method does require a complete distribution and open tails must be extrapolated to some arbitrary grain size. Another assumption is that each class to be measured is normally distributed, which is generally false and can distort the results. Other pitfalls include the fact that the mode, or for that matter, any polymodality, cannot be determined using the method of moments. Graphical methods can be very time consuming if performed manually, however, computerized methods are available to retrieve points from the graphs. This method is considered to be statistically less valid than the method of moments. Results only approximate those obtained by the method of moments since 5-8% of the distribution is missing. Of course, this could also be a benefit if the distribution has open tails or if the tails contain enough uncertainty, due to measurement errors, to justify removal.

2.2 Image Acquisition

Thin-section analysis is conducted with the use of transmitted light microscopy. Samples are prepared by first impregnating them with a dyed epoxy, usually blue, but other colors can also be used. A thin-section is then cut from the sample, mounted to a glass slide, then ground and polished to a thickness of 30um. The thin section sample can then be examined with a microscope using two different types of lighting, each of which produces unique optical cues [Gribble and Hall, 1992].

2.2.1 Plane Polarized Light

Placing a polarizer in the path of light causes it to vibrate in single plane. This is then referred to a plane polarized light (PPL). Using this type of light, the field of view appears bright and generally white, depending on the particular type of rock being analyzed. The following features can be observed using PPL [Gribble and Hall, 1992].

- Transparent phases appear as white since they allow the majority of light to pass through. These phases could be minerals, glass or liquid (depending on the presence of dye in the epoxy).
- Absorbing phases appear black since they absorb most of the light. Typically opaque material or ore minerals.
- Grain boundaries, cleavage traces and micro-fractures appear as thin black lines. This is due to the fact that light is scattered and refracted at these locations.
- Fluid inclusions can sometimes be observed as small irregular rounded areas within the minerals. They are often found grouped together or in zones.
- Holes, fractures, and places where the rock section is missing appear white (depending on the presence of dye in the epoxy).
- Artefacts, such as bubbles in the mounting medium, may appear either in areas filled with epoxy or superimposed on the minerals themselves. Other possible artefacts include preparation materials, such as grinding grit, which may accumulate in fractures or at the edge of a section. This material will appear as a dark colored dust.

Systematic methods have been developed to described minerals in thin-section using plane-polarized and crossed-polarized transmitted light sources. The following describes the methods implemented when using plane-polarized light [Gribble and Hall, 1992].

Color – the colors of minerals can vary significantly from the clear minerals such as quartz and feldspars to the opaque ores that appear black. The color of a particular mineral is dependent on the wavelength of light that it allows to pass.

Pleochorism – refers to the change of color experienced by some minerals when the microscope stage is rotated. The mineral will exhibit two extremes separated by 90 degrees of rotation. This is due to the orientation of the wavelengths of light.

Habit – defines the shape exhibited by a specific mineral within the rock type. Euhedral minerals will show well-defined crystal faces while anhedral minerals have no distinguishable crystal facies. Typically, euhedral minerals can only be properly developed when there is nothing to interfere with the crystal growth. Minerals that are elongate in one direction are termed prismatic; those that are needle-like are known as acicular. When the minerals resemble fibers, they are termed fibrous. Finally, flat thin crystals are described as being platy or tabular.

Cleavage – refers to the planes along which a mineral can be cleaved or split. These planes are straight, parallel and evenly spaced in the mineral. Quartz does not exhibit cleavage. Poorly developed cleavages are referred to as partings and are generally straight and parallel but not evenly spaced.

Alteration – The result of chemical reactions between minerals and water or CO2. These reactions can be so advanced that the original crystal is completely replaced by the new

stable mineral phase(s). In such a case, when the shape of the original mineral is maintained, it is considered to be pseudomorphed by the new mineral.

2.2.2 Crossed Polarized Light

To achieve crossed polarized light, the use of an analyzer is required. The analyzer is similar to the polarizer except it causes the light to vibrate at right angles to the polarizer. When light from the polarizer passes through the analyzer, the effect is a dark image, since light, that is vibrating perpendicular, is not passed through. Images, viewed in crossed polarized light (XPL), have a different set of terminology that is used to describe the minerals. Generally, the use of XPL is required to properly identify all of the image constituents found within a thin-section petrographic image [Gribble and Hall, 1992].

2. 3 Review of Petrographic Image Analysis

Initial work in the area of petrographic image analysis (PIA) was focused on extracting information about the pore geometries. This early work provided most of the justification and motivation for the use of PIA and laid the groundwork for more advanced analysis. The basis for this work is the stereological concept known as the Delesse Principle [Delesse, 1847], which suggests that representative quantitative three-dimensional data can be derived from two-dimensional petrographic images. To test this theory, porosity data from PIA was compared to values obtained from standard volumetric methods and the results show good agreement between the two methods [Ruzyla, 1986]. In addition, multiple thin sections were taken from a single core plug and compared to each other and to volumetric methods to determine if a single thin section image could be representative

of the entire plug. Again, the results from these tests suggest that a single thin section could indeed represent the properties of the plug. The real test of the benefits and possible uses of PIA is whether it is capable of predicting other reservoir properties, such as permeability, which are more applicable to economic considerations. Permeability is considered to be the most important property of reservoir rock since it dictates how the reservoir fluids move throughout the reservoir. Gies and McGovern [1993] use pore measurements to develop a simulated capillary pressure curve, which in turn can be used to estimate parameters such as permeability, effective porosity, and irreducible water saturation. In other cases, the identification of distinct pore types allowed the definition of rock types, thereby leading to explanations of reservoir heterogeneities and permeability [Davies, 1990; Bowers et al., 1994; James, 1995]. In other work, diagenetic material was removed from samples using image-processing techniques. This allowed for the quantification of diagenetic effects in terms of permeability reduction [Mowers and Budd, 1996]. Much of this work involving pore characterization and its relationship to permeability is based around fundamentals developed by Ehrich et al. (1991a, 1991b) and McCreesh et al. (1991). Other workers stress the integration of imaging techniques to give a complete characterization of rock structure [Tomutsa et al., 1990; Radaelli et al., 1998]. This includes Scanning Electron Microscopy, X-Ray Computed Tomography and Magnetic Resonance Imaging. Methods have also been developed to estimate grain-size distributions and mineralogy [Clelland and Fens, 1991; Starkey and Samantaray, 1994; Francus, 1998], but these measures are not limited to sandstone. Samples of igneous rock [Thomson, 1996; Ross et al., 2001], coal [Lester et al., 1993], and unconsolidated sand [van den Berg et al., 2002] are also examined. In addition, grain size distributions are

important in other polycrystalline materials and their analysis is similar to PIA [Diogenes et al., 2003; Kraatz et al., 2003; Mahadevan and Casasent, 2003].

2.3.1 Petrographic Image Processing and Analysis

Methods of image acquisition vary, depending on the application. For identification of pore space, thin-sections can be viewed with an optical microscope using plane-polarized light. For additional grain detail and determination of mineralogy, crossed-polarized light must be used. A scanning electron microscope (SEM) set in the backscattered electron (BSE) mode can also be used to identify pore space and when an electron dispersive x-ray analyzer (EDX) is added it can be used to identify mineralogy. Cathodoluminescence is used to investigate diagenetic effects and to examine components of polycrystalline grains.

Images acquired with an optical microscope are known to suffer from the Holmes effect [Crabtree et al., 1984]. This phenomenon, also known as the shelving effect, is due to the fact that the samples are a finite thickness and therefore the edges of grains tend to fade away. BSE images are considered to be a true two-dimensional image and therefore grain boundaries are more distinct. Other pitfalls of transmitted light microscopy include the fact that image brightness depends on sample thickness, which may vary, and matrix material tends to be blurry. Despite these drawbacks, thin-section images require less preparation time as compared to BSE images and therefore remain one of the cheapest and quickest methods for making gross measurements on single images [Francus, 1998; Ross et al., 2001].

Each of the above acquisition methods will produce images that can be processed and analyzed automatically, however, image processing and analysis was not the focus of the early work and very few details are typically given. This is mainly due to the fact that only the pore space was needed and for methods such as BSE the acquired image was distinctly bimodal and only required thresholding to provide the segmentation of pore space from the remaining image. Eventually, focus shifted to additional analysis of the structure including grain characteristics and mineralology. This type of analysis proved to be difficult as polycrystalline grains were over-segmented and mineral overgrowths were counted as individual grains [Clelland and Fens, 1991]. Also, many grains remain fused together and require operator interaction to separate them. These types of errors have a significant impact on measures of grain size distribution

One method used to eliminate touching grains modifies the means of image acquisition. It captures numerous sequential images under varying conditions of illumination and then accumulates them to produce a synoptic image [Starkey and Samantaray, 1994]. This acquisition method is meant to replicate the operator interaction that is required for polarized light microscopy with the final image containing more information than any single image. Despite these efforts manual interaction was still required to separate touching grains.

One of the first formal methods developed for petrographic analysis involved filtering to reduce noise, image segmentation via operator determined gray-level thresholding, image editing to fill holes within the objects, and manual separation of touching grains [Francus, 1998]. This method has since been expanded to include

correction of non-uniform lighting, linear contrast stretching to enhance contrast and maximum likelihood classifier to identify grain material [van de Berg et al., 2002].

Efforts have been put forth to fully automate the separation of individual grains in order to facilitate quicker processing of thin-section images. Erosion-dilation cycles were presented [Russ, 1995] but are known to distort the grains original shape. Watershed methods were also presented [Russ, 1998; Zhao, 2000; van den Berg et al., 2002] but they tend to give a large number of erroneous results. An attempt was made to reduce the number of false watershed separations by using a modified watershed transformation [Zhao, 2000] and while this method had some success, many errors remained. A variation of the above methods, known as the digital cutting method (DCM), was based on connecting characteristic contact wedges that form between touching grains [van den Berg et al., 2002]. The overall results of this method showed only moderate improvement as compared to the watershed method.

2.3.2 General Image Processing and Analysis Concepts

One needs to look no further than the seemingly endless lists of journals dealing with image processing and its related topics to gain an appreciation for the shear volume of work carried out in this area over the last three decades. It has grown in leaps and bounds spurred on by the desire to produce systems capable of replicating (in some form or another) the human visual system. The motivation for this is often blurred with some researchers doing theoretical work in the area and others using basic well-known methods to develop ad hoc systems suitable only for a specific purpose or use. In recent years, more effort has been focused on developing general methods capable of understanding

real world images [Casadei and Mitter 1998]. It is pointed out that much of this work on general segmentation and understanding is based the perceptual grouping concepts of: similarity, proximity, continuity, symmetry, parallelism, closure and familiarity [Malik et al., 2001]. In other words, algorithms are geared more towards image understanding and are increasingly robust, requiring less *a priori* knowledge of the image composition. Proper development of these algorithms requires an intimate knowledge of image processing and its related disciplines. This is not an easy task since these disciplines include image processing, image understanding, scene analysis, machine vision, digital signal processing, artificial intelligence and others.

Petrographic images can vary considerably since they are the result of natural processes. Francus [1998] indicates that image analysis techniques must be adapted to the optical characteristics of each set of samples and a generally accepted standard of analysis has yet to be developed. Many formal image analysis methods were investigated for this work but in the end an algorithm was developed based on straightforward and well-known concepts. Excellent overviews of image processing concepts can be found in several sources [Russ, 1995; Gonzalez and Woods, 2002]. Background that is specific to microscopy can be found in [Russ, 1990] while a review of concepts related to thin-section analysis was conducted in [Zhao, 2000].

Chapter 3

Manual Analysis of Data Set

The purpose of this chapter is to investigate the characteristics of the images available for analysis and to select images for use as a data set for further testing. Also, the process of manually segmented these images is examined and criteria are established for identifying and classifying quartz grains.

3.1 Image Characterization

Thin-section images were acquired from five separate cores spanning approximately 240m of a single well in the Hibernia formation [Hibernia]. Three different images are available for each interval corresponding to three levels of magnification. A macroview is included at a magnification of X7.5 to show the overall structure of the specimen. Next, a low magnification image is taken at X63 to show the grain and pore space characteristics. Finally, a high magnification of X125 is used to observe the smallest features and details. A total of 72 images are available for this study; 36 low magnification images known as group A and 36 high magnification images labeled group B. The main focus is on group A since at this magnification there should be a sufficient number of grains to give an accurate size distribution. These images are 880 by 649 pixels and represent a field of view measuring 1.80 mm x 1.33 mm. A variety of image features can be expected since many of the pictures were taken to indicate

heterogeneities. Also, the images span several reservoir structures, thereby, increasing the likelihood of showing a range of features. While this provides a good data set for testing the robustness of the final image processing routine, it also significantly increases the difficulties associated with designing such a routine. An attempt is made to address the majority of the most important and frequently occurring features present.

The thin-section images used in this study are siliciclastic and contain whole and fragmented quartz, mostly of irregular shapes, due in part to the rounding and fragmentation that occurs during the depositional process. The quartz may be monocrystalline and appear as a single distinct grain, or polycrystalline and appear to be composed of several smaller grains. In addition, structures may be altered by any cementation that may have taken place during the diagenetic process, resulting in complex grain boundary interactions.

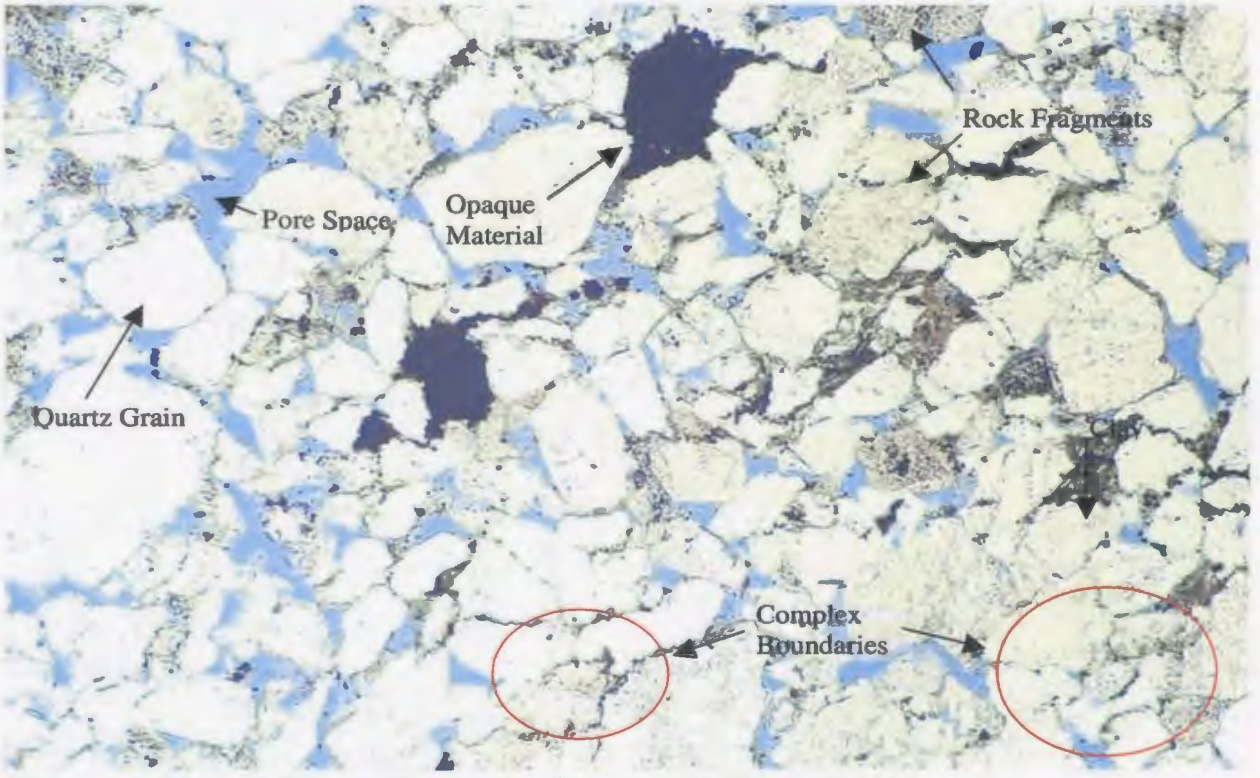

Figure 3.1: Typical thin-section image demonstrating complexity.

Other materials present in the image include clay, rock fragments, and opaque material such as organic detritus and certain heavy minerals. Figure 3.1 shows an example of a typical thin-section image. The blue represents the pore space in the image and results from a blue dyed epoxy used to hold the specimen together during preparation. Another notable feature produced by the preparation process is what in known as the 'shelving effect', which causes the grains to fade away in places where boundaries are shared with pore space. It is also cited as a partial cause for some of the complex grain boundary interactions often observed in thin-section [Van den Berg et al., 2002]. The term 'typical', used to describe the above image, is a bit of a misnomer since the thin-section images vary significantly throughout the data set. The only guaranteed substance is the quartz that seems to appear in every one of the images. All other constituents vary significantly throughout the set as will be demonstrated in the next section.

3.1.1 Summary of Image Characteristics

A qualitative study of the test images was conducted to gain insight into the features present and the manner in which they vary from image to image. These features can be grouped into three main categories: image defects, diagenetic material, and grain characteristics.

The image defects of most concern include the non-uniform lighting, the overall appearance in terms of the image brightness and the extent of the shelving effects. These defects are a result of the image acquisition methods and must be corrected before proceeding with any image analysis. Image processing methods are available to correct non-uniform lighting and poor brightness. The shelving effects, however, require a

completely different acquisition method that would allow for a true two-dimensional image and cannot be corrected using image-processing techniques.

The diagenetic material includes opaque material, quartz overgrowths, clay and silt material as well as evidence of dissolved grains. These features are important in the interpretation of the reservoir geology and give indications of the depositional environment through an understanding of the diagenetic history of the formation. In terms of image segmentation, the diagenetic material complicates the process by precipitating between grains and causing them to appear joined. It also can form around detrital grains with a boundary separating the original grain and the diagenetic material; however, this boundary is not always visible in plane light. The amount of diagenetic material varies significantly throughout these images.

Grain characteristics, such as the average grain size, presence of uniform textured grains, and the number of dirty grains are also noted. Many of the high texture or dirty grains are not considered to be quartz, however, in some cases polycrystalline quartz grains show complex internal structure and therefore can also be classified as high textured.

A very significant range in properties and features are present in the A group images. Taking into consideration the above features, the 36 images were rated according to the overall image quality. Only two images were identified as having good overall image characteristics while 15 images were rated as moderate and 19 were identified as having poor overall quality. The good images, as illustrated in Figure 3.2, are unambiguous with respect to the image components and it is expected that such a low number of these images would be present.

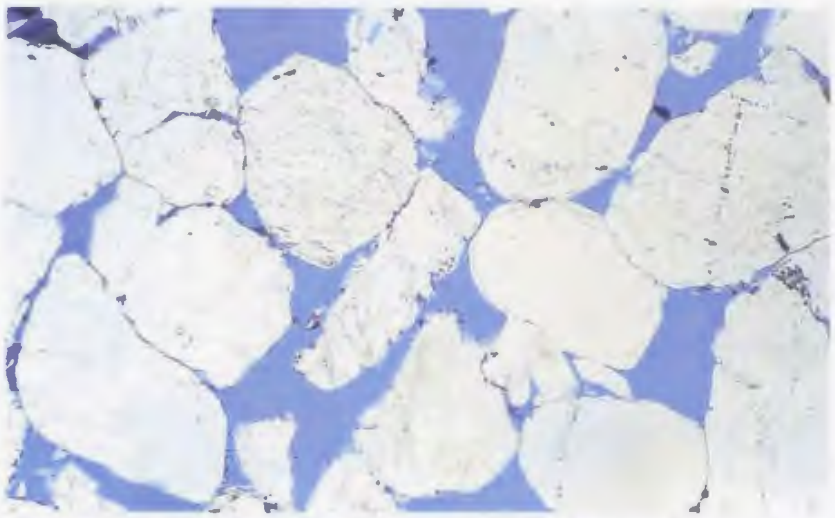

Figure 3.2: Good quality image with distinct grains.

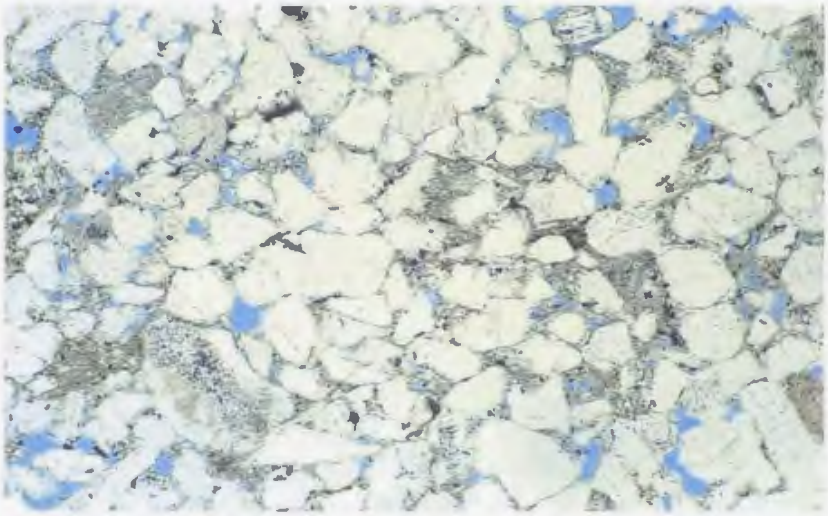

Figure 3.3: Moderate quality image with majority distinct grains.

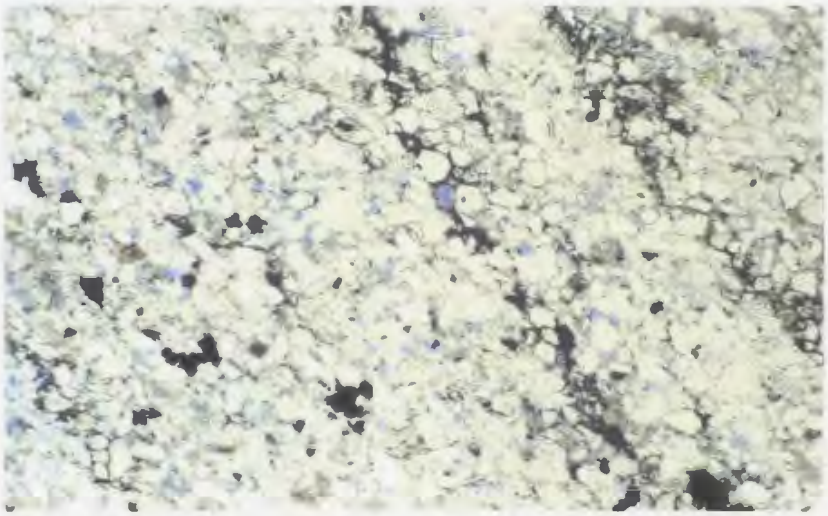

Figure 3.4: Poor quality image with few distinct grains.

The images that are rated as moderate, shown in Figure 3.3, appear to contain a majority of easily identifiable grains, however, some complex regions are present. Finally, poor images, shown in Figure 3.4, contain a variety of features that make it difficult to identify image constituents including: lighting effects, diagenetic alteration, high texture areas, and small grains.

3.2 Manual Segmentation

Three moderate quality group A images were selected and manually segmented by a senior reservoir geologist with HMDC. While not an expert in petrographic analysis, his knowledge of reservoir geology provided a sound basis for the manual segmentation. The three thin section images are shown in Figures 3.5A, 3.6A, and 3.7A along with their corresponding manually segmented images, Figures 3.5B, 3.6B, and 3.7B.

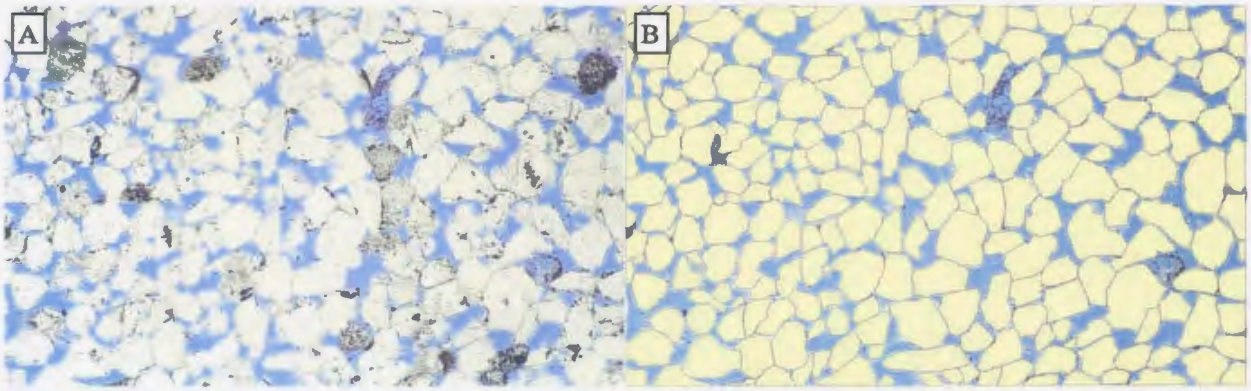

Figure 3.5: Image 1 A) Original image. B) Manually segmented image.

These three images were re-segmented by an inexperienced operator with the primary goal of investigating the way in which the images are interpreted with limited *a priori* knowledge of petrology. A secondary goal was to study the methods used by an experienced geologist.

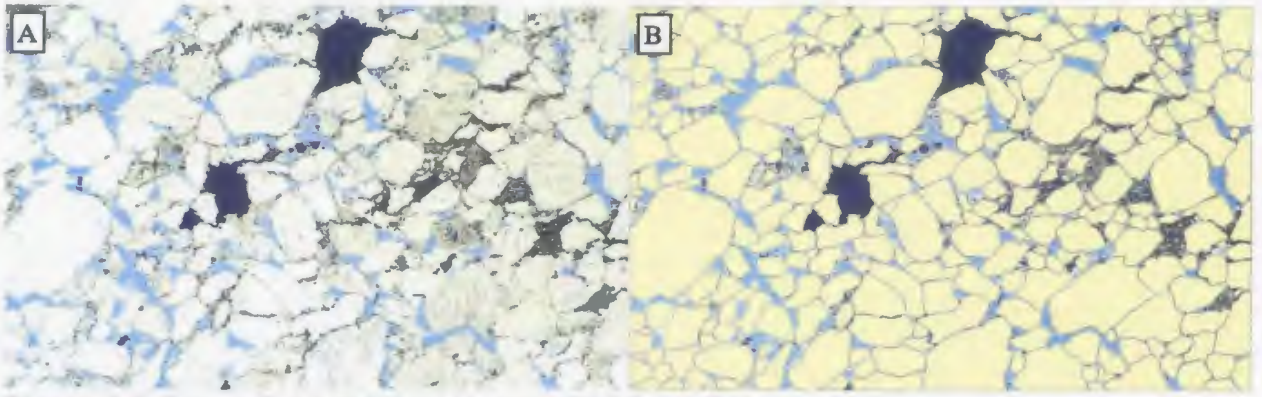

Figure 3.6: Image 2 A) Original image. B) Manually segmented image.

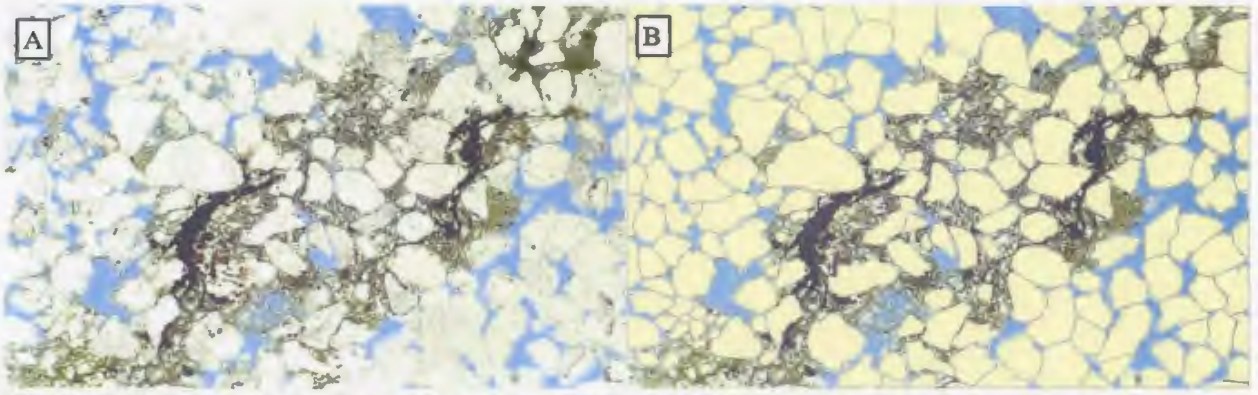

Figure 3.7: Image 3 A) Original image. B) Manually segmented image.

The results, Table 3.1, show significant differences in interpretation and in each case the number of segmented grains identified by the inexperienced operator is greater. This suggests that many visual cues are present to prompt the operator to over-segment the image.

		Average Area	Normalized Std Dev of Area	Number of Grains	Percent Total Image Area
Image 1	Operator 1	1509	0.67	232	61.3
	Operator 2	1181	0.87	305	63.1
Image 2	Operator 1	1254	1.45	282	61.9
	Operator 2	1258	1.37	317	69.8
Image 3	Operator 1	1140	1.13	211	42.1
	Operator 2	1276	0.99	242	54.1

Table 3.1: For each of the 3 images different operators provide different interpretations.

Further study of these manually segmented images suggested that no formal guidelines were employed in the segmentation, i.e. rules that were applied to some grains were not applied to other similar grains. Such results can be expected, since human interpretation can vary depending on a number of factors [Russ, 1990].

In addition to the images shown in Figures 3.5-3.7, 11 other group A images were selected and manually segmented to give 14 images in total. Out of these images two are of good quality, eight are rated as moderate, and the remaining four are considered to be poor. A selection of 14 group B images was also manually segmented and used for comparison. While manually segmenting these images, a classification scheme was developed to identify problematic grains. This scheme is the focus of the following section.

3.2.1 Grain Types

Classifying the grains based on their complexity is an important concept. A fundamental goal of automating the thin-section analysis is to increase precision by limiting the effects of operator bias. The majority of these variations occur when interpreting the complex features of certain types of grains. Disregarding these grains for textual analysis would therefore eliminate some of the operator bias. Also, without a set methodology for segmenting these grains, it would be impossible to quantify the performance of the developed segmentation algorithms. It is theorized that the textual properties of the image can be accurately characterized by determining the features of certain 'types' of grains. This simplifies the process to some extent, and mimics the manual methods currently used to perform textural analysis. For example, the laboratory that completed

the analysis on this set of images used 100 monocrystalline quartz grains from each sample to provide grain size analysis [El-Dein et al., 1984]. Also, if comparative charts are used to determine sorting (or other features) the operator would not mentally segment all grains, just the ones needed to give an indication of the overall trend. Finally, and most importantly, it is not possible to fully segment the entire image based on a one field of view. As discussed in Chapter 2, the analysis of these types of thin-section images requires constant operator interaction to accurately identify all image constituents.

Type 1 Grains

Type 1 grains, which are the easiest to identify, have complete boundaries separating them from surrounding grains and pore space. In addition, type 1 grains are uniform with respect to their intensity values, they contain no significant holes or defects, and they have a fairly regular shape as shown in Figure 3.8.

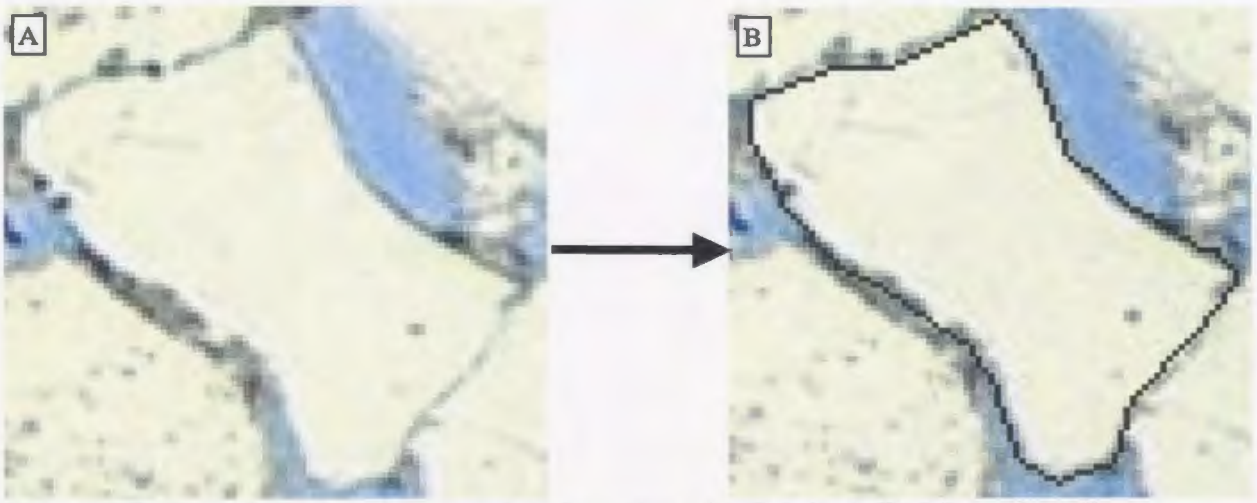

Figure 3.8: Type 1 Grain with complete boundary A) Original. B) Manual outline.

In geological terms, type 1 grains are detrital unaltered non-overlapping mono-crystalline quartz grains as seen using plane polarized light. These grains are either matrix

supported, and therefore not in contact with any other grains, or are in point contact. Some type 1 grains may share concavo-convex or sutured contacts provided there is a complete boundary, although, this is seldom the case. It is expected that an inexperienced user can identify type 1 grains with 100% certainty.

Type 2 Grains

Type 2 grains are defined as those grains that are not completely separated from each other but posses properties that allow for easy visual segmentation from the surrounding background. Reasons for incomplete boundaries include shelving effects, moderate quartz overgrowths, low degrees of compaction (pressure solution), and poor lighting conditions. Type 2 grains can possess either point, concavo-convex, or sutured contacts with incomplete boundaries as shown in Figure 3.9. Overly complex sutured contacts with missing segments will cause the grains to be classified as type 3 grains. It is expected that an inexperienced operator would be able to correctly identify 70-80% of type 2 grains and with some basic knowledge of petrographic analysis it would be possible to identify the remaining 20-30%.

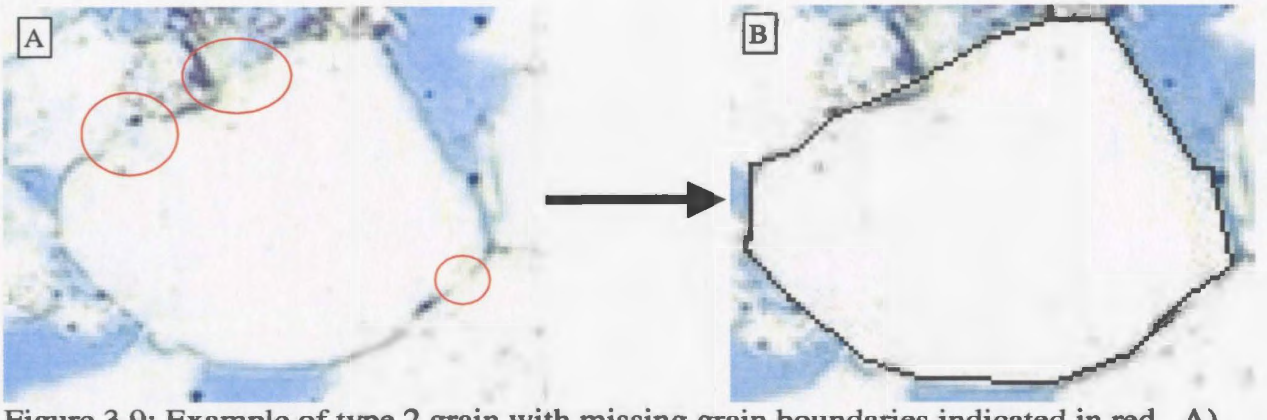

Figure 3.9: Example of type 2 grain with missing grain boundaries indicated in red. A) Original. B) Manual outline.

Type 3 Grains

Type 3 grains are those that cannot be segmented manually, either because there is no satisfactory interpretation or because there is more than one interpretation. Generally speaking, these regions of grains are either composed of polycrystalline quartz and/or some combination of other materials or they have been diagenetically altered. These regions lack the clear boundaries and distinct textural features that would enable an operator, either skilled or unskilled, to properly complete the segmentation. A typical grain region that is classified as type 3 is shown in Figure 3.10. Two significantly different interpretations of the same field of view are shown which illustrates the difficultly associated with these grains. The segmentation requires analysis of extremely subtle features and any interpretation would be a guess.

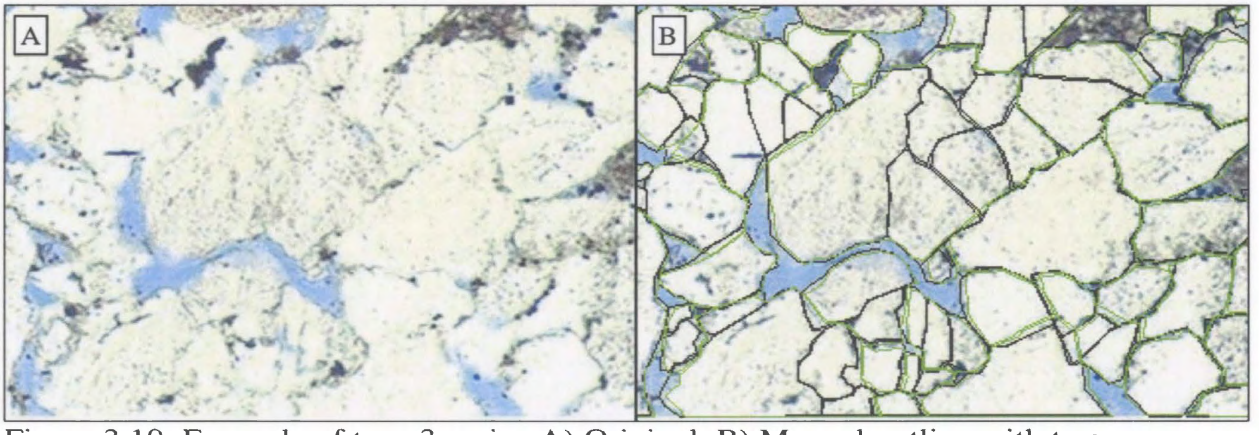

Figure 3.10: Example of type 3 grains A) Original. B) Manual outline with two interpretations. Green is an experienced operator, black is inexperienced.

The only true way to interpret a polycrystalline grain is to view it under crossed polarized light. Figures 3.11A and B [El-Dein et al., 1984] demonstrates this concept. Both images show the same field of view, the first utilizes plane-polarized light and the second results from crossed-polarized light. The small arrows indicate mono-crystalline quartz while the large arrows indicate polycrystalline quartz.

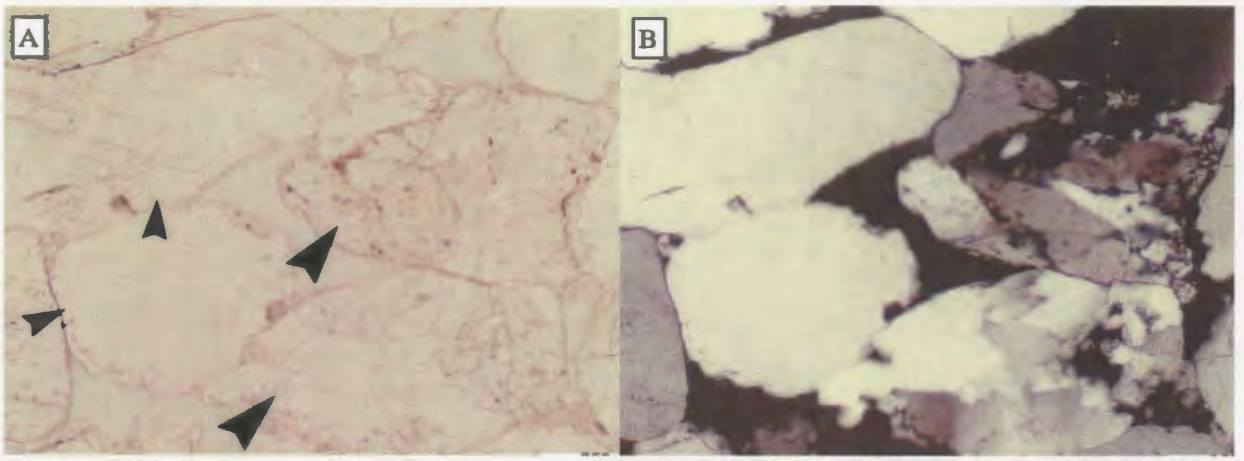

Figure 3.11: Polycrystalline Quartz – same field of view. A) Plane-polarized light. B) Crossed-polarized light. Large arrows indicate polycrystalline quartz while small arrows indicate monocrystalline quartz [El-Dein et al., 1984].

Polycrystalline grains present significant challenges to the development of an automated segmentation routine. As can be seen in Figure 3.11, part of the polycrystalline grain appears to be several smaller grains stuck together, whereas, the remaining portion of the grain appears to be mono-crystalline (until viewed under crossed polarized light).

Borderline Grains – Uniform Texture

Grains with uniform textural characteristics, as shown in Figure 3.12A, have the potential to fall within the guidelines of type 2 grains provided they could be easily identified. However, these types of grains are generally not quartz grains, but instead are composed of rock fragments, chert, feldspars, or other materials. It is not the purpose of this algorithm to accurately identify and segment non-quartz grains; therefore, it is suggested that these grains should be classified as type 3 grains. Another variation of a grain with uniform (or semi-uniform) texture is shown in Figure 3.12B. The texture is quite different from the grain shown in Figure 3.12A and the dark parallel lines could be

cleavage traces indicating that it may be a feldspar grain (since quartz does not exhibit cleavage).

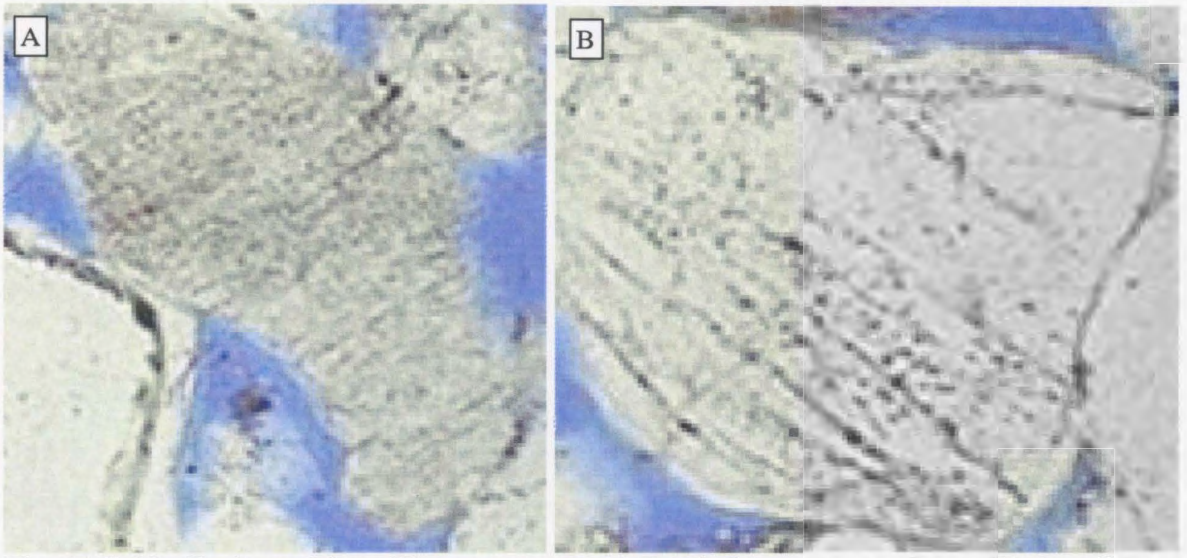

Figure 3.12: Grains with uniform texture. A) Rock fragment or chert. B) Possibly feldspar, fluid inclusions, or vacuoles.

Technically, this grain should not be included as a type 2 grain since it may not be quartz. Such a grain could cause problems for an automated segmentation routine since the dark parallel lines fall within the same range of intensity values as the grain boundaries and, therefore, keeping or eliminating one set of lines would require keeping or eliminating the other. It may be beneficial to consider these to be type 3 grains.

Borderline Grains – Diagenetically Altered

All grains have been affected, to some extent, by diagenetic processes since it is these processes that change the sand into sandstone. This can manifest itself in several forms including fracturing, overgrowths, cementation and dissolution. In Figure 3.13A quartz overgrowths are clearly identified by the presence of dust on the original grain surface. It is clear that these grains should be classified as type 3 since they have the potential to be

over-segmented (i.e. overgrowths will be removed). In some cases, when these layers of dust are not present the overgrowths can still be easily identified due to their euhedral shape, as in Figure 3.13B, but typically some other method is needed to quartz overgrowths. Grains without a distinct layer of dust should therefore be classified as type 2.

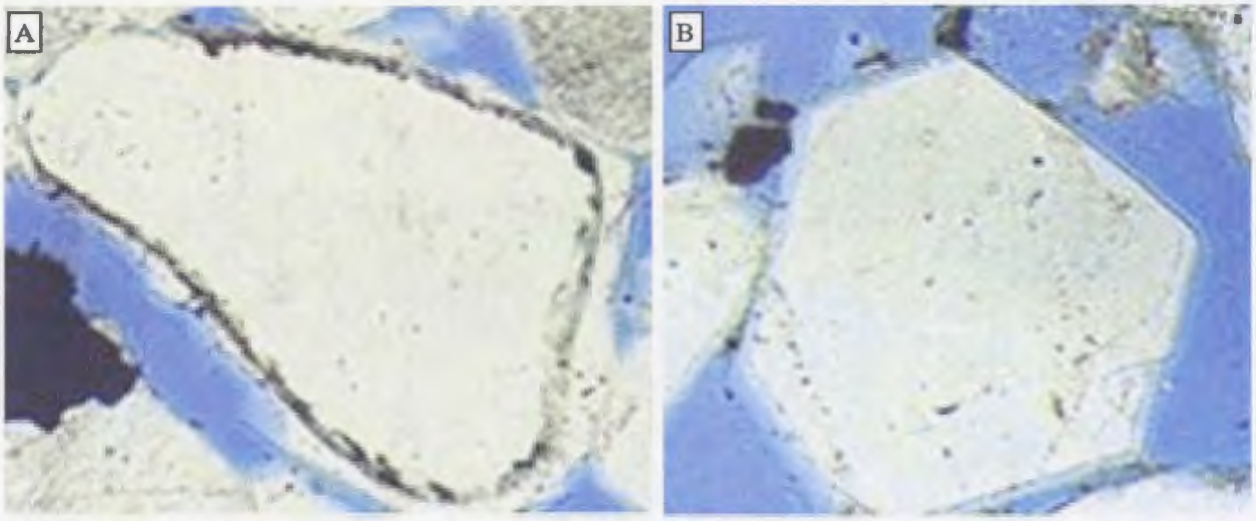

Figure 3.13: Distinct quartz overgrowths. A) Visible due to dust on original grain. B) Visible due to euhedral shape.

Another example of complex diagenetic alteration is shown in Figure 3.14A. Figure 3.14B shows an experts interpretation of this grain. It is clear that this type of grain should be classified as a type 3 grain, since it requires piecing together many smaller distinct grain regions and this type of interpretation is beyond the scope of this thesis.

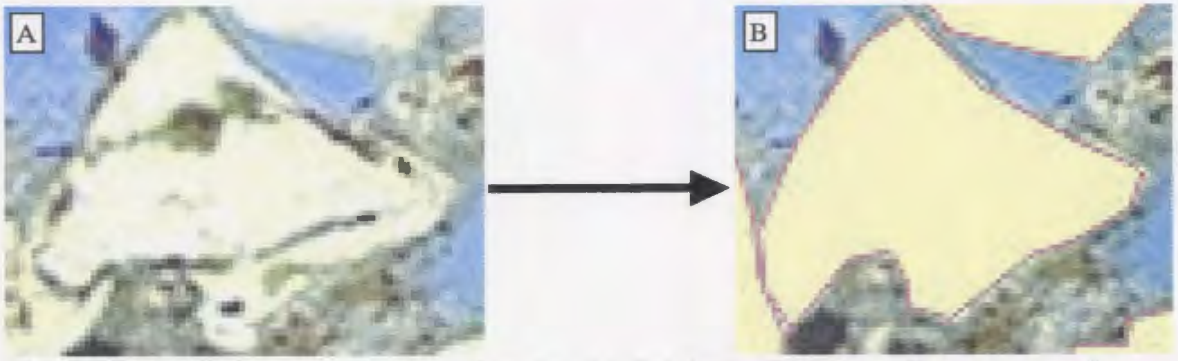

Figure 3.14: Diagenetically altered grain. A) Original. B) Manually segmented.

Figure 3.15A shows a large partially dissolved polycrystalline grain. The dissolution has started to take place along the boundaries, causing a confusing situation for segmentation. The combination of edges and pore indentations is the criteria used to segment many other grains but in this case the complexity of the interior grain boundaries and the presence of a well-defined outer boundary outweighs the other information and suggest that this grain should not be segmented. Technically, this grain should be classified as a type 2 grain, however, there is a tendency to label these grains type 3 since there is a degree of uncertainty associated with the identification.

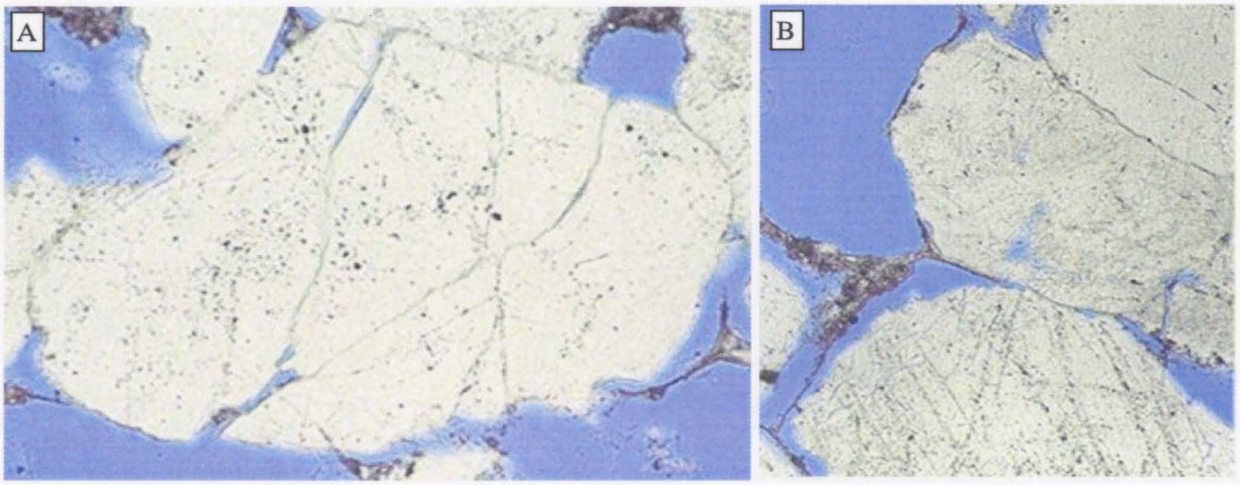

Figure 3.15: Dissolution of quartz grains. A) Polycrystalline grain with dissolution along internal boundaries. B) Dissolution around external boundary and within grain. Clay shows location of original boundary.

Figure 3.15B also appears to be partially dissolved with blue stain in the interior of the grain. This type of grain is also difficult to classify since the porosity suggests segmentation while the overall shape indicates a complete grain. This particular grain can be classified as type 2 but variations where the porosity is more pronounced should be classified as type 3.

The final two types of diagenetically altered grains are: fractured, Figure 3.16A and cemented Figure 3.16B. The fractured grain is an extreme case and most other examples show lesser degrees of fracturing.

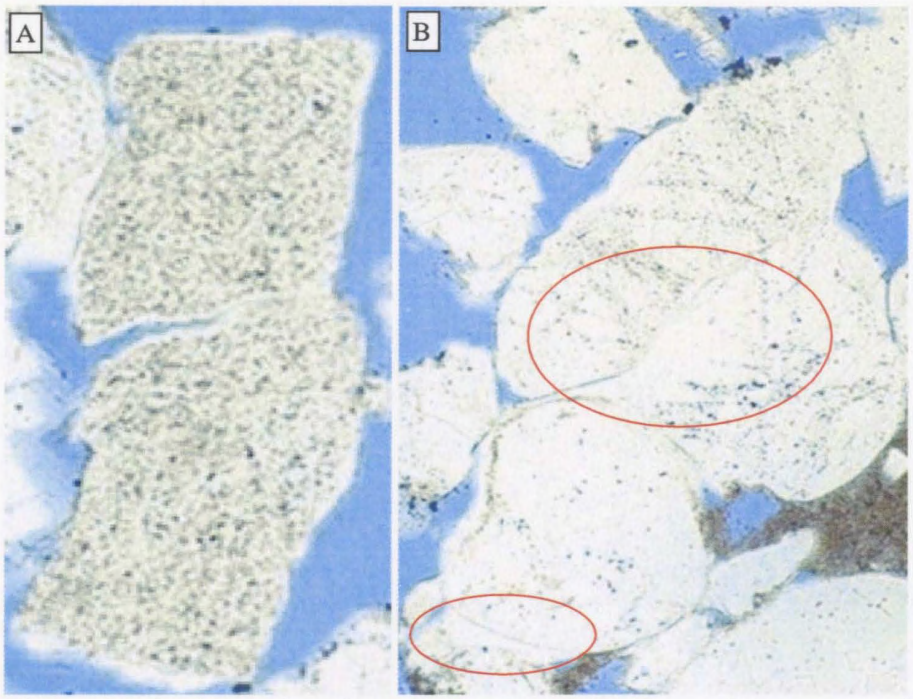

Figure 3.16: Altered grains. A) Fractured. B) Excessive cementation.

This type of grain is considered to be borderline since the porosity indentation suggests splitting but the overall texture and shape suggests it may be considered as a complete grain. It is unclear which interpretation is correct and therefore this is classified as a type 3 grain. The cemented grains are the result of excessive quartz overgrowths and in some cases it is clear where the original grains end and where the cement begins. In this particular case, the boundaries are unclear therefore prompting this entire grain region to be classified as type 3.

3.2.2 Summary of Grain Type Data

To investigate the effects of neglecting type 3 grains, images in the data set were re-segmented and the type 1, type 2, and type 3 grains were identified. Figure 3.17 shows an example where type 3 grains are illustrated in dark blue while the type 1 and type 2 grains are pale yellow. Six measurements were recorded for the manually segmented grains; they include: area, roughness, compactness, elongation, breadth and width. These features are calculated using the MIL Blob analysis feature [Matrox, 1999].

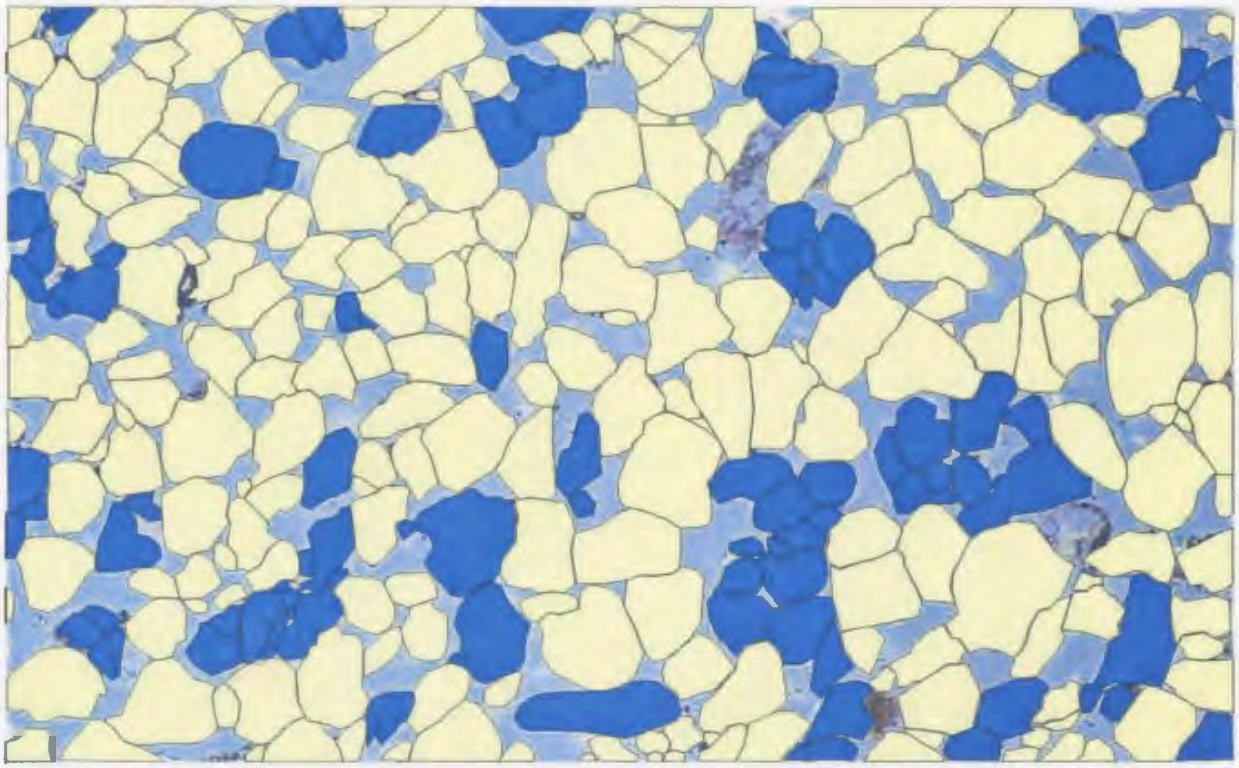

Figure 3.17: Manually segmented image - blue grains classified as type 3.

Figure 3.18 shows the average area of individual grains for each image while Figure 3.19 shows the standard deviation of the area.

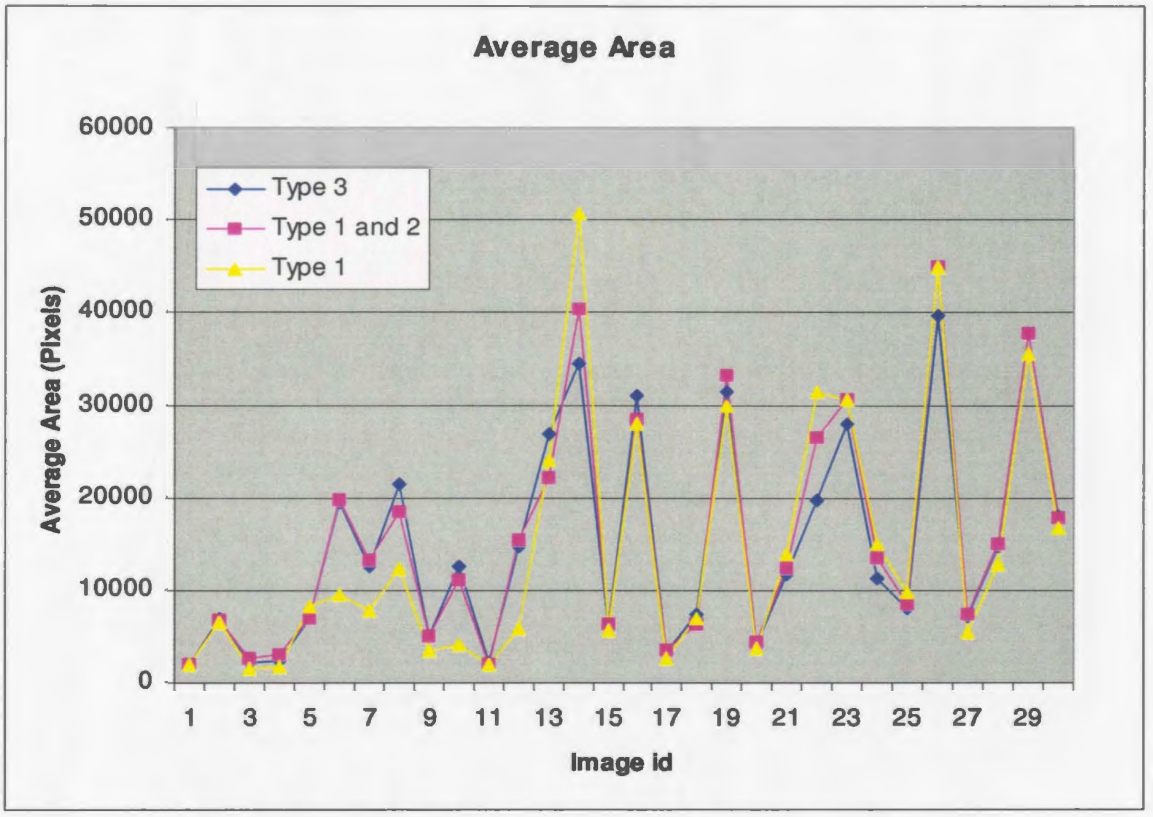

Figure 3.18: Average area calculated for the different grain types.

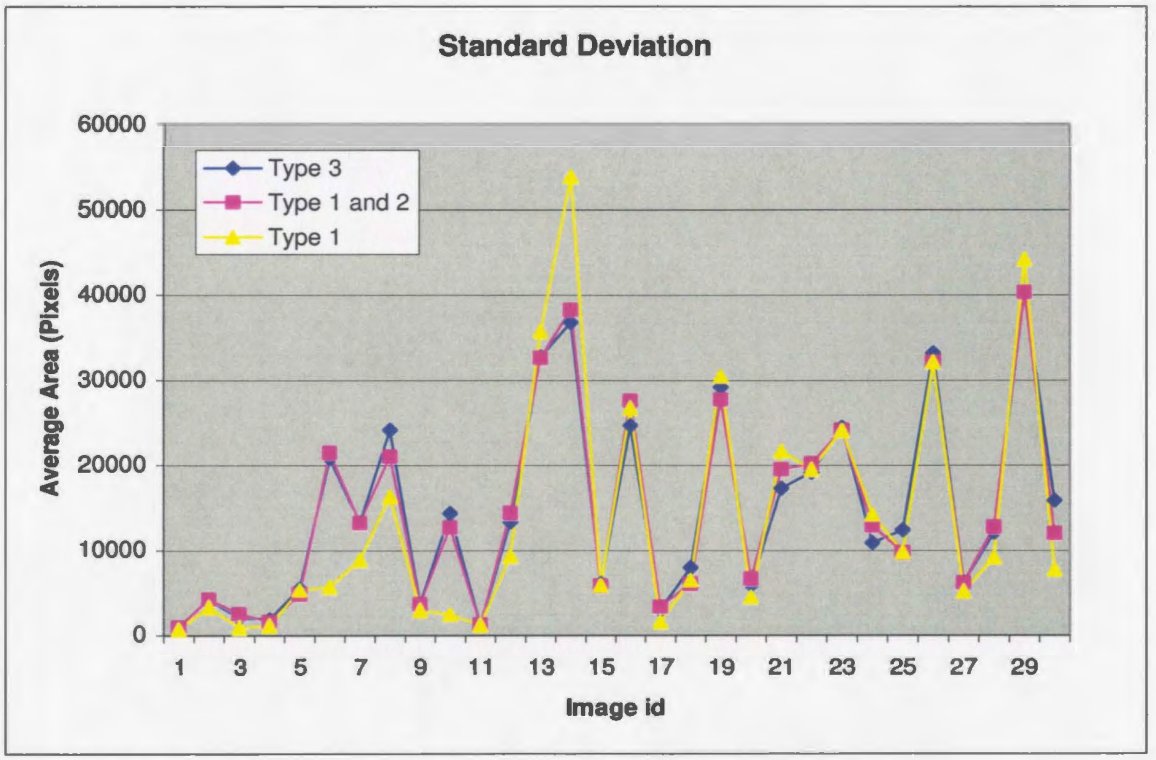

Figure 3.19: Standard deviation of the area calculated for each of the grain types

The figures show that the average area and standard deviation values for type 1&2 together are closely related to the values for all grains. The type 1 grains alone seem to correlate well with all grains, suggesting that only a small number of grains are required to give a good estimate of the overall distribution. Pearson product moment correlation coefficients were calculated for all features and their standard deviations and the results are shown in Table 3.2.

Pearson	Area		Breadth		Compactness		Elongation		Length		Roughness	
	Ave	SD	Ave	SD	Ave	SD	Ave	SD	Ave	SD	Ave	SD
1&2 with All	0.97	0.93	0.95	0.94	0.94	0.92	0.94	0.9	0.97	0.96	0.87	0.58
1 with All	0.89	0.86	0.84	0.82	0.72	0.79	0.73	0.78	0.91	0.86	0.37	0.1

Table 3.2: Pearson correlation coefficients for all images.

The coefficients show a strong positive relationship between type 1&2 grains and all grains. Table 3.3 shows the correlation coefficients calculated for just the group A images.

Pearson	Area		Breadth		Compactness		Elongation		Length		Roughness	
	Ave	SD	Ave	SD	Ave	SD	Ave	SD	Ave	SD	Ave	SD
1&2 with All	0.99	0.99	0.93	0.99	0.98	0.96	0.98	0.95	0.98	1.00	0.87	0.56
1 with All	0.98	0.99	0.90	0.90	0.73	0.66	0.76	0.66	0.96	0.98	0.77	0.47

Table 3.3: Pearson correlation coefficients for A group images

When considered separately, the group A images show an even stronger correlation. This is most likely related to the fact that the number of grains found in the higher magnification group B images is on average much lower than those found in group A. The above data shows that the type 3 grains can be left out of the segmentation process without negatively impact the overall characterization of the image.

3.3 Grain Shape Characterization

Grain characterization is not only required as part of the textural analysis process, it is also needed for the segmentation algorithm. An important step in the algorithm is to identify grains once they have been separated from one another. This is not an easy task due to the complex nature of the rock-forming process. As mentioned in the previous section, six measurements were recorded for the manually segmented grains; they include: area, roughness, compactness, elongation, breadth and width. The area, roughness and compactness were selected for grain characterization.

3.3.1 Area Considerations

For textural analysis, only grains within the 4 – 0.03 mm diameter range are considered significant. The lower end of this range corresponds to the boundary between coarse silt and medium silt. This works out to be approximately 180 pixels for the A group, which is a fairly small object in these images, so consideration was given to moving the lower cutoff up one step on the phi scale to 0.0625 (1/16) mm diameter. From initial study of the images, it seems that this is a logical cutoff since objects below this limit could correspond to diagenetic overgrowths or other partial grains and therefore should not be consider complete grains. The goal now is to detect and measure sand sized grains with all grains in the silt range being classified as non-grains. Analysis of the manual images shows that only about 4.3% of the total grain area is composed of silt sized grains. Although the silt grains are not significant in terms of area, they do make up nearly 37% of the total number of grains in the A group images. If the accuracy of the segmentation routine is based upon the number of correct grains found, the uncertainty in identifying

silt-sized grains could significantly impact this value, possibly to such a degree that it becomes meaningless. So silt-sized grains will be identified but will not be taken into account when determining the accuracy of the algorithm. It is also worth noting that the petrographic report completed by El-Dein et al. [1984] considers the silt-sized grains to be part of the matrix material.

3.3.2 Roundness and Compactness

Roughness and compactness were selected to measure the angularity and sphericity of the grains. To investigate the effectiveness of these measures, they were applied to Figure 3.20, a standard reference chart taken from Lewis and McConchie [1994b].

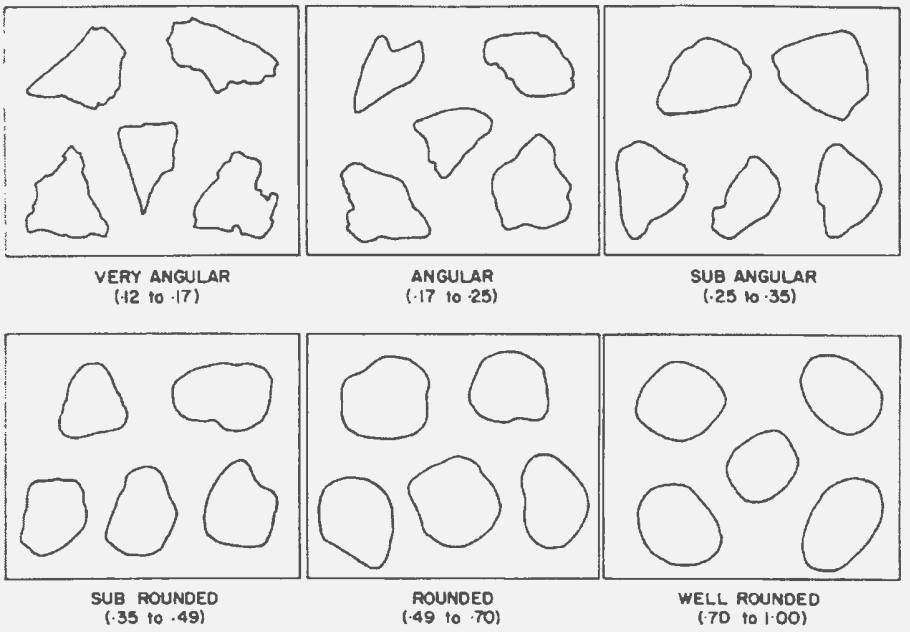

Figure 3.20: Rounding Silhouette [Lewis and McConchie, 1994b].

Table 3.4 shows the values obtained by applying the roughness and compactness measures to the grains in Figure 3.20. Compactness values increase as the grains become less rounded, which is expected since lower values of compactness indicate objects that are close to circular. The same trend is found with the roughness values. They increase

with increasing angularity, which is expected since higher roughness values indicate a greater degree of angularity. These features are measured for the manually segmented group A images. About 45% of the grains have compactness values within the sub rounded to sub angular range, which compares well with the petrographic analysis that classified all samples as sub-rounded to sub-angular [El-Dein et al., 1984].

	Compactness	Roughness	Average Compactness	Average Roughness
Very Angular	1.94	1.15	2.04	1.16
	1.89	1.14		
	2.01	1.09		
	2.11	1.15		
	2.25	1.26		
Angular	1.62	1.14	1.82	1.15
	2.27	1.18		
	1.81	1.14		
	1.60	1.14		
	1.81	1.14		
Sub Angular	1.45	1.12	1.62	1.13
	1.53	1.13		
	1.65	1.13		
	1.60	1.12		
	1.85	1.16		
Sub Rounded	1.40	1.10	1.44	1.10
	1.48	1.10		
	1.44	1.11		
	1.49	1.12		
	1.38	1.09		
Rounded	1.32	1.09	1.36	1.10
	1.30	1.09		
	1.42	1.10		
	1.38	1.12		
	1.38	1.09		
Well Rounded	1.34	1.11	1.35	1.11
	1.32	1.11		
	1.37	1.12		
	1.41	1.11		
	1.32	1.11		

Table 3.4: Compactness and Roughness measures

Only 3.3% of the compactness values are above 2.25, which is the maximum value in Table 3.4. This indicates a possible cutoff value to testing whether or not an object is a grain. The roughness data obtained from the manual images does not correspond as well. Over 90% of the grains have roughness values in the well-rounded to sub-rounded range. This can be explained by considering the straight lines used to estimate the grain boundaries in the manual segmentation process. Much of the detail is lost, thereby, reducing the roughness of the edges and producing lower roughness values. Roughness is still considered to be a good estimate of the grains angularity but some trial and error will be required to determine the cutoff value for determining grain criteria.

3.3.3 Sorting

Methods for sorting were introduced earlier but both the method of moments and the graphical methods have their drawbacks. For this application, sorting is not of utmost importance but some measure is needed to give an indication of how well the algorithm works compared to the manual segmentation. Standard deviation was the preferred method of sorting measurement implemented in previous work on this data set [Zhao, 2000]. To obtain a value similar to the Phi scale, the standard deviation is normalized by dividing by the mean value to obtain a relative standard deviation. Table 3.5 shows the relationship between Phi standard deviation, relative standard deviation, and the verbal scale.

PHI standard deviation	Relative standard deviation	Verbal scale
0.35	0.2377	Very well sorted
0.5	0.4017	Well sorted
1	0.5542	Moderately sorted
2	0.9989	Poorly sorted

Table 3.5: Sorting Measures

Chapter 4

Development of Segmentation Methods

It takes highly skilled operators to properly analyze a thin-section petrographic image. They draw upon knowledge of how the rocks are formed, (sedimentary petrology), and the manner in which the images are acquired (optical microscopy). The intuitive feel that they develop for how to properly segment the images can be linked to the perceptual grouping concepts introduced in Chapter 2. It seems that the challenge is to find the right combination of image-processing and analysis routines that are able to mimic these concepts. Identification of type 1 grains should be straightforward and is the focus of the primary segmentation routine. From studying the thin section images, it is apparent that in most cases type 2 can be segmented using a combination of corner, line, and texture information as demonstrated in Figure 4.1.

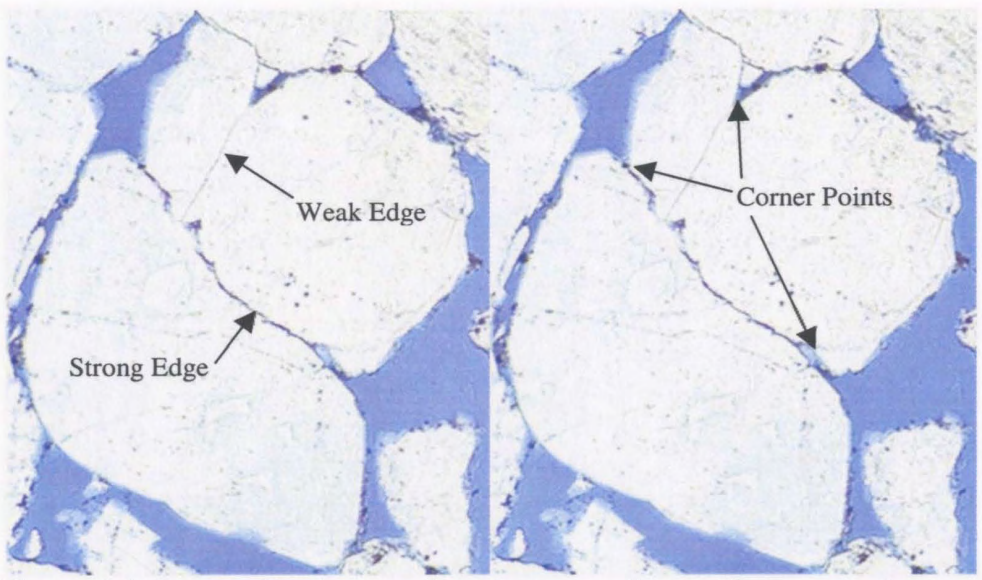

Figure 4.1: Image showing weak and strong edges and corner points used by an expert to segment thin section images.

Hence, efforts are focused on connecting partial and weak line segments to corners and other, more prominent, line sections in an attempt to replicate the perceptual process used by an expert. This is the job of the secondary segmentation routine. After all of the segmentation has taken place, the grain reconstruction routine makes an attempt to reconstruct some of the over-segmented grains. The full algorithm is illustrated in Figure 4.2. In addition to the primary and secondary segmentation routines, it includes preprocessing methods as well as a routine to classify the image constituents. Once grains and other material have been classified, statistics are gathered and output to a text file. The preprocessing and primary segmentation are inter-related and, therefore, are discussed together.

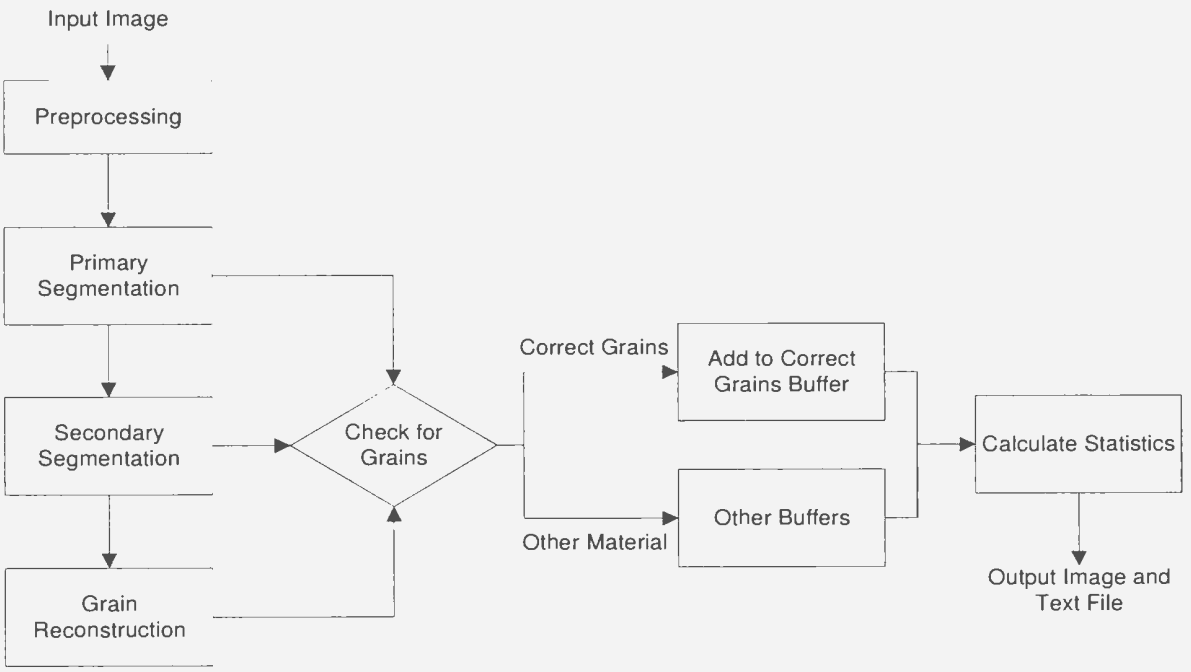

Figure 4.2: Algorithm overview

4.1 Primary Segmentation

The goal of the primary segmentation routine is to produce an image containing only type 2 and type 3 grains, which can then be used as input for the secondary segmentation routine. This process involves removing the clay, silt or opaque material from the image along with any type 1 grains. The methods used here are based mostly around gray scale thresholding techniques with some filtering used to reduce noise and blob analysis [Matrox, 1999] in order to classify and measure image features. The red channel is chosen for further processing since it has the highest contrast between grains and pore space [Zhao, 2000].

4.1.1 Filtering

Several types of filters exist for noise reduction; the mean, median and Gaussian filters are among the most popular. Starkey and Samantaray [1994] conducted a comparison of filters for petrographic analysis; however, due to the fact that these types of images can vary significantly, it is still necessary to investigate filtering for each specific application. Previous work with this data set concluded that the median filter provided a good means of noise removal for these images [Zhao, 2000]. While the median filter is an effective means of noise removal, it also has a tendency to reduce the visibility of grain boundaries, as shown in Figure 4.3, and in the case of weak boundaries it results in their removal. Also, repeated applications of the filter tend to reduce the presence of larger spots and other grain features that are important to the segmentation process. Francus [1998] used a hybrid median filter that eliminates noise but preserves edges. This

approach was considered here but in many cases it is just as important to eliminate edges as it is to eliminate noise.

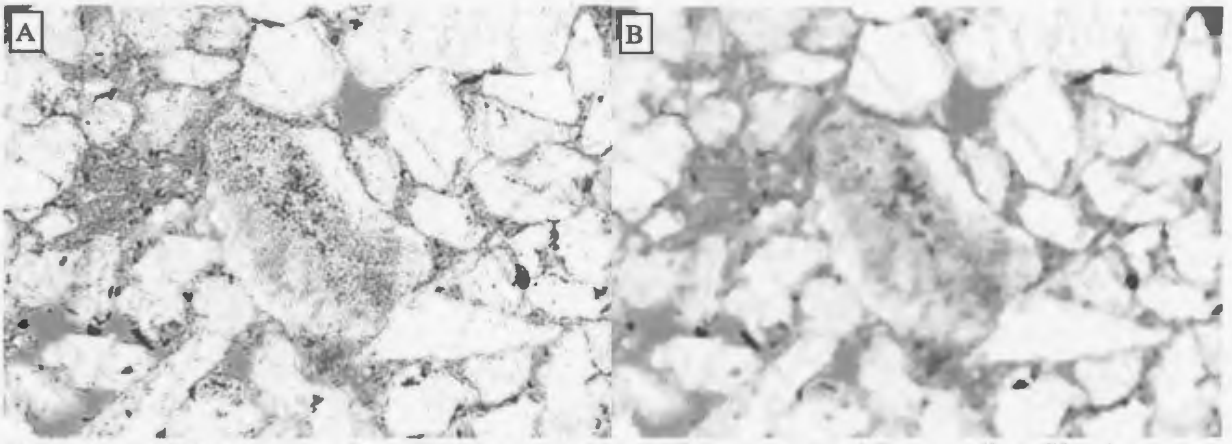

Figure 4.3: Effects of median filtering. A) Original image. B) After median filtering.

Further consideration of the trade-off between noise reduction and edge preservation is provide in later sections. It is clear that at least one application of the median filter is needed, but it is unclear at this point whether or not additional applications are required.

4.1.2 Thresholding

Thresholds were investigated for all of the 36 A group images to determine the range in intensity values for each of the important image features. These intensity ranges were compared to the histograms for each of the images to determine whether or not there is a relationship between these values and peaks in the histogram. Examples of histograms for the red channel are shown in Figure 4.4.

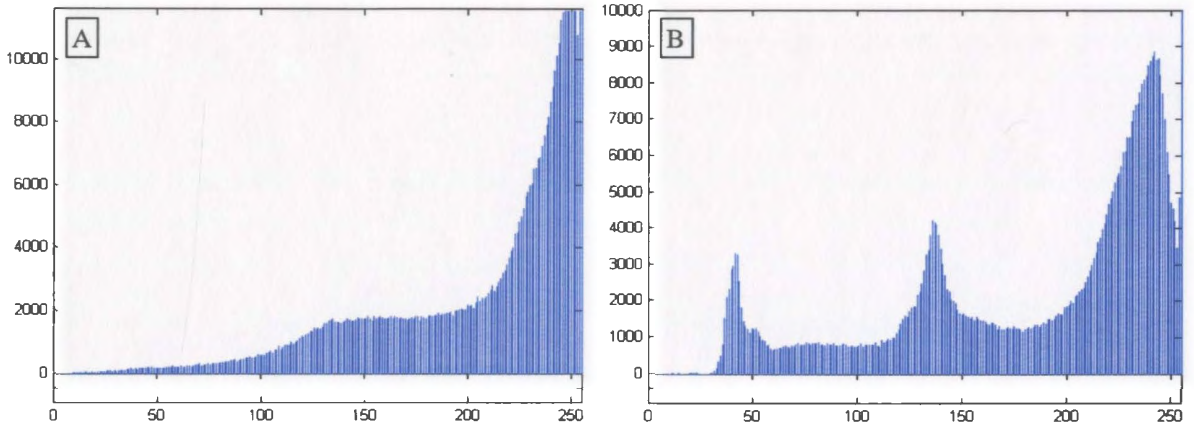

Figure 4.4: Histograms showing one large peak corresponding to quartz grains (A) and three peaks representing opaque material, pore space and quartz grains (B).

Histograms vary from image to image but in most cases the intensity ranges correspond to specific image features as indicated in Table 4.1.

	Opaque Material	Pore Space	Shelving & Dirty Grains	Normal Grains	Clean Grains
Intensity Range	80-130	130-170	170-210	210-230	230-240

Table 4.1: Intensity ranges for main image features.

For the opaque material, a single threshold is required and all pixels below that value can be considered opaque material. In this case, the threshold can vary from 80-130, depending on the image. In terms of area, opaque material can be very significant in some images and therefore has the potential to provide valuable information for characterization. The thresholding and blob analysis can provide a quick measure for determining the amount and type of opaque material. Large circular blobs may be authigenic pyrite or siderite while large elongated blobs are likely plant detritus [El-Dein, 1984].

For the pore space, shelving effects, and dirty grains, a high and low threshold is required and the features are more or less composed of the intensity values within the given range. The dirty grains appear much darker than other grains and show either

uniform or random texture features. These grains are not likely to be quartz grains and as discussed in Chapter 3, they are classified as types 3 grains. Therefore, no further efforts will be made to segment and measure the dirty grains using thresholding methods.

Finally, for the normal and clean grains, the threshold is single-valued and all pixels with intensity values above the threshold are considered to be grain pixels. The distinction made between normal and clean grains is to indicate that some grains are brighter and have uniform intensity values falling within a smaller range. This effect is partially due to non-uniform lighting conditions and partially due to the types of grains present. It was noted that many clean monocrystalline grains are present in these samples [El-Dein, 1984].

The normal grain intensity range contains the majority of quartz grains. Raising the threshold value results in better separation of grains in the center of the image but at the expense of loosing grains on the right boundary. Several solutions to this problem were investigated including subtracting a blank image, fitting a background function and rank leveling [Russ, 1995]. The approach taken here is based on the adaptive thresholding method developed by Zhao [2000] in which the image is divided into four equal regions and thresholds are then determined automatically from each region's histogram. It was found that while this worked well for some images, better overall results could be obtained by using 32 sub-images as opposed to just four. The method used to determine the thresholds is based upon examination of the histogram. The intensity value with the greatest number of pixels associated with it is selected as the maximum histogram value. These peaks are only considered if they are greater than 180 so, for example, peaks due to pore space are not considered for thresholding. Once the

peaks are determined, two different correction factors are applied to the maximum histogram value to obtain a threshold value. For sub-images with maximum histogram values above 230, a high correction factor is subtracted from the intensity value to give the threshold. Sub-images with values below 230 have a low correction factor subtracted to give the threshold. For the preliminary work the high and low correction factors are set at 30 and 20 respectively. Each sub-image is then binarized using its histogram-based threshold. This method reduces the threshold values used for the sides of the image and produces an improved binary image as can be seen in Figure 4.5; note that the grain material is white.

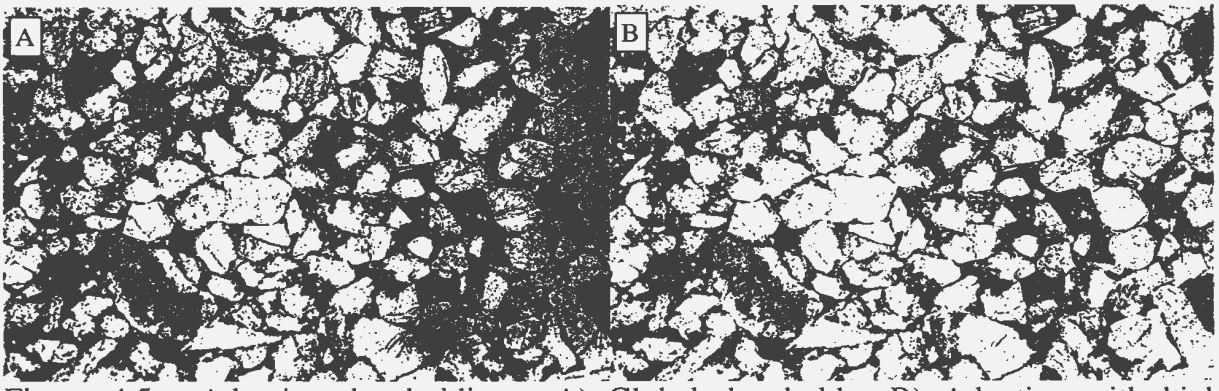

Figure 4.5: Adaptive thresholding. A) Global threshold. B) Adaptive with high correction factor set to 30 and low set to 20.

It is important to note that grain detection/segmentation is sensitive to the threshold correction factors and the lower the factor the better separated the grains tend to be. The tradeoff is an overall reduction of grain area, which could significantly alter the shape of the grains.

4.1.3 Type 1 Grain Detection

After initial thresholding is performed, a number of type 1 grains have been segmented and the focus is shifted to detecting these grains. Objects with area values less than 50

pixels are labeled as dirt to indicate a lack of interest in this size range and to distinguish them from the medium and coarse-grained silt that lies in the 50-732 pixel range. The dirt size range encompasses the fine and very fine silt grains. Clay is not distinguishable in grain form and appears as a brown matrix that contains silt sized grains dispersed throughout. The blob analysis package [Matrox, 1999] is used to eliminate objects (blobs) based on their area. These objects, once removed from the image, are placed in separate buffers and kept for future reference. Figure 4.6 shows the results of removing the dirt-sized objects.

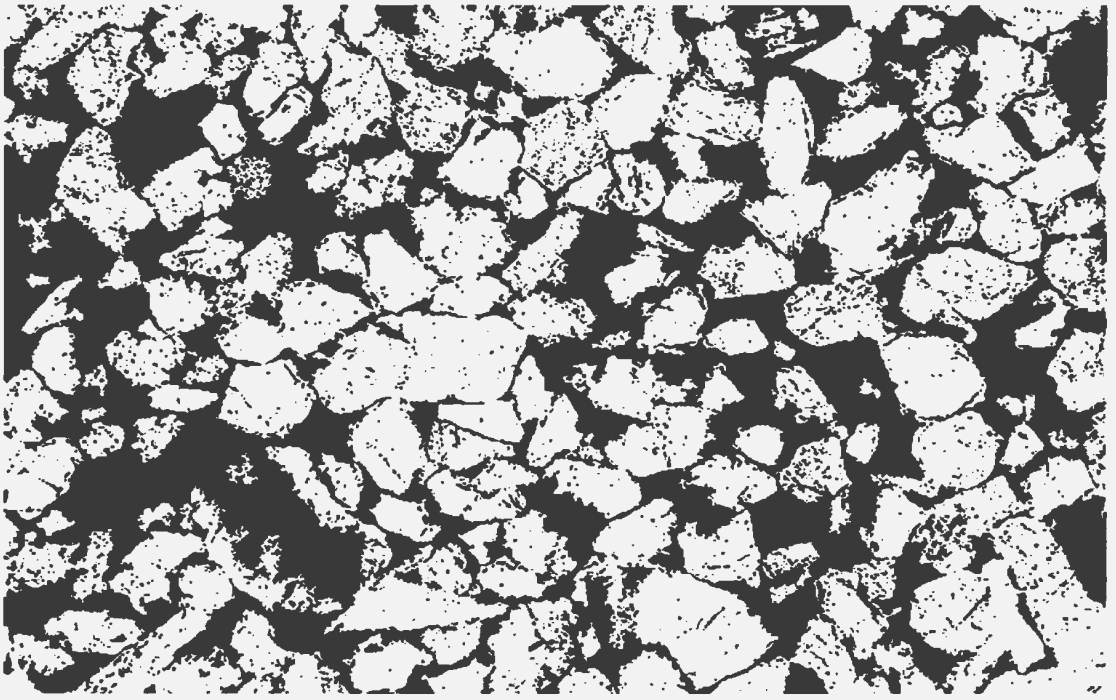

Figure 4.6: Removal of dirt-sized objects.

The small black blobs within the grains (fluid inclusions or clay on original boundaries) can also be removed using the blob analysis package. This presents an alternative to excessive filtering of the grayscale image and allows a means of selecting features of interest. Excluding blobs that have low compactness values (circular shapes) leaves only straight blobs that are more likely to be partial grain boundaries. Some constraint on the

area is still required since too many blobs remain that correspond to "dirt". The area threshold is set to 50 for this particular image, however, images that have different sized grains would require different threshold values. For example, segmentation of smaller grains may benefit from retaining blobs less than 50 pixels in area. Figure 4.7 shows the result of eliminating interior "dirt", based solely on area and Figure 4.8 indicates, in red, the benefits of using the compactness criteria.

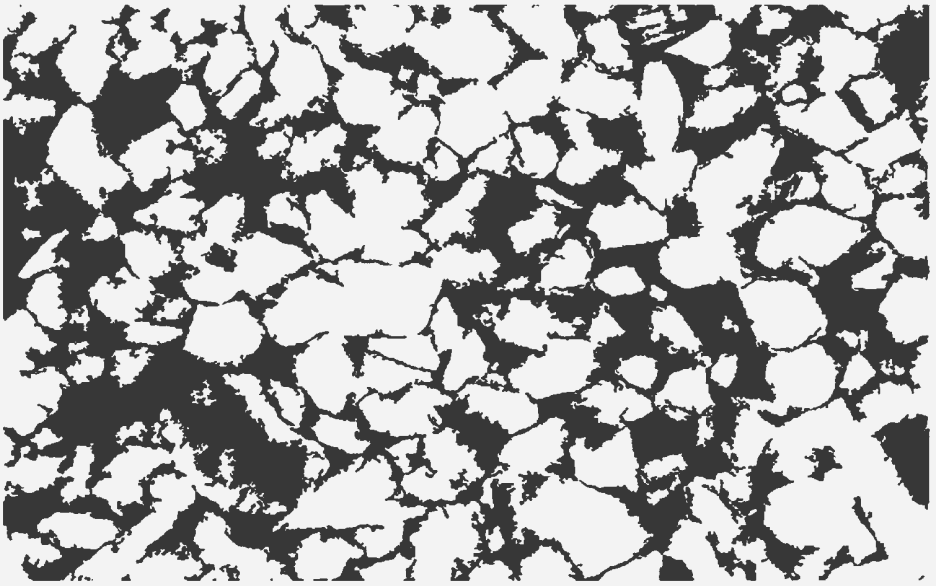

Figure 4.7: Interior dirt removed based on area constraints.

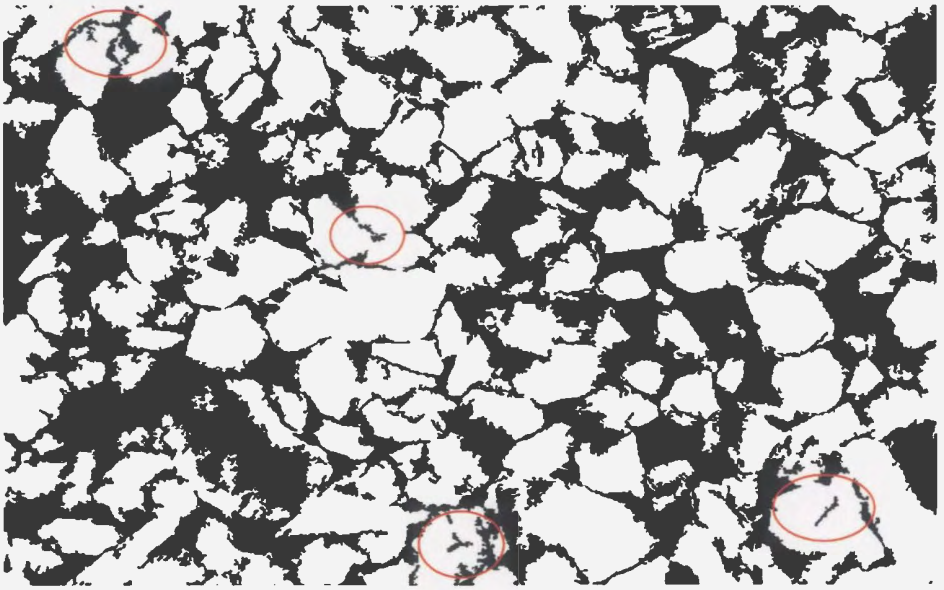

Figure 4.8: Interior dirt removed based on area and compactness constraints. Red indicates features that are important for segmentation.

Once the image has been cleaned using the blob analysis, the search for grains can begin. Roughness and compactness are the criteria used to detect the grains and an object is only considered to be a grain if it meets both criteria. A number of possible roughness and compactness values were tested based on the discussion in Chapter 3. In the end, it was found that many correct grains were being rejected based on these criteria. The reason for this is linked to the rough grain edges produced by the thresholding process. Applying the median filter to the grayscale image will tend to smooth these boundaries, however, this has side effects and it would be better to have a method of boundary smoothing that is independent of the grayscale image. The proposed solution involves applying morphological closing operations to each grain before calculating the roughness and compactness values. The idea is that the closing will fill in the small gaps along the boundary without significantly altering the shape of the grain. The closing operation leaves grains with large indentations unaltered and, therefore, they can be left for the next segmentation routine. This process allows roughness and compactness values to be set in accordance with the analysis provided in Chapter 3. Objects meeting the grain criteria are removed from the image and placed in a buffer, to be measured.

4.1.4 Detecting Other Type 1 Grains

Type 1 grains are defined in Chapter 3 to have complete boundaries when viewed as an unaltered color image. After filtering and thresholding, these grains may be connected. These connections often consist of just a few pixels and simple methods can be used to separate these type 1 grains from the rest of the grain fabric.

Morphological Operations

Erosion dilation cycles have been proposed to separate touching objects and while this method has had some success it does so at the cost of distorting the grain shape. This is not acceptable in this case since information about the shape of the grain is required. However, this method could be used conservatively so that the shapes are not significantly changed. For example, one application of morphological opening would result in the separation of objects that are connected with only one or two pixels.

Edge Detectors

Edge detectors can also be used to segment type 1 grains. Boundaries that may not have been detected by the thresholding routine can be detected using the edge operators since they respond to absolute differences in intensity values. Several edge-finding routines were applied to the red channel image including the Sobel and Canny detectors. While the Canny edge detector is often considered to be the optimal routine, its added complexity was not required for this application and the Sobel operator did a good job of finding edges. The Sobel operator is applied to the red channel and the result is a grayscale image where the intensity values represent the strength of the edges. Applying a threshold to this image is equivalent to selecting edges of a particular strength; higher intensity values correspond to stronger edges. The binary image produced by the thresholding is then skeletonized and blob analysis is conducted to remove small closed loops that are not needed for the segmentation. The focus is on detecting more significant edges. The clean skeleton image is then inverted and overlaid on the binary grain image and any newly segmented grains are analyzed to determine their agreement

with the grain criteria. Newly classified material will be placed in their respective buffers.

4.1.5 Estimating High Texture Material

High texture regions, composed mainly of clay material, do not fall into a specific intensity range and they are well suited for textural analysis. The most important consideration for this problem is that the clay material varies with respect to intensity in all directions. Edge detectors are good for measuring intensity gradients and, therefore, can be used to detect high texture areas. For this application, the Sobel edge detector was applied to the red channel. For clay regions, a high response is recorded; this is then thresholded and the holes are filled to produce complete objects that represent the clay material. This image is then measured to provide an estimate of the amount of clay material in the image. Regions that do not have intensity values varying in all directions are not detected using this method. Therefore line segments, such as those shown in Figure 3.12B, are not measured but dirty grains that are likely rock fragments, Figure 3.12A, are included.

4.2 Secondary Segmentation

Secondary segmentation is based on approximating missing grain boundaries of type 2 grains. Several methods are available to estimate the grain boundaries. Work has been carried out to connect characteristic indentations, or contact wedges, that indicate the presence of grain boundaries [van den Berg et al., 2002]. This concept was examined and it became clear that it would only be successful on a specific type of grain contact.

Because of the diagenetic alteration found in this data set grain contacts seldom meet the above criteria.

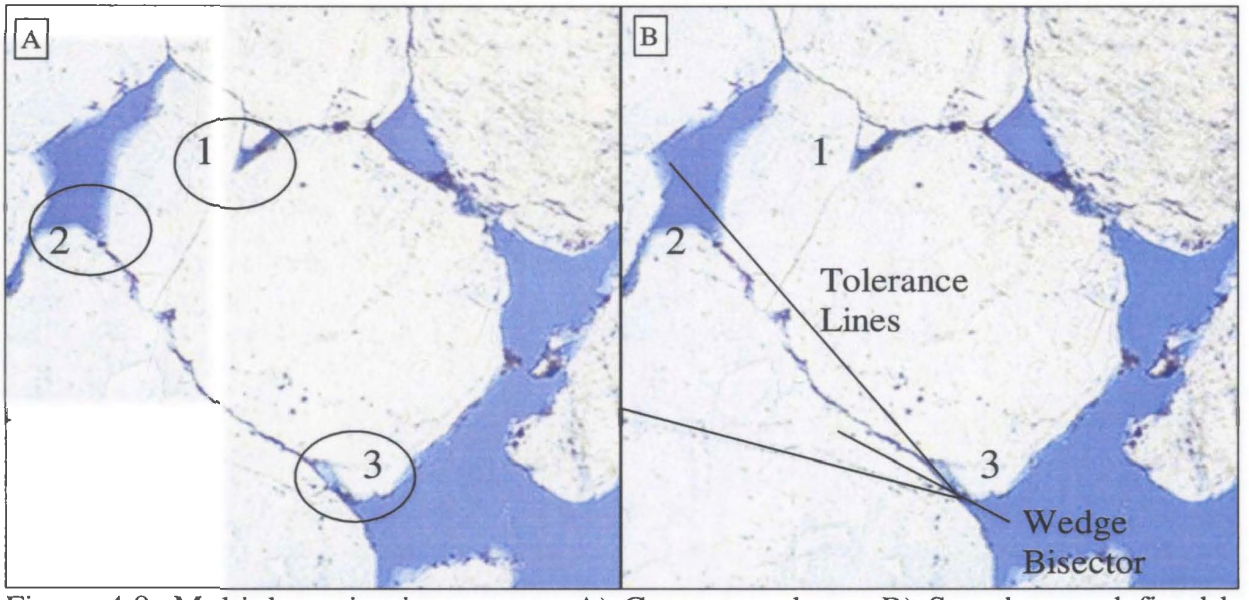

Figure 4.9: Multiple grains in contact. A) Contact wedges. B) Search area defined by tolerance lines.

This is largely due to the fact that two or more grains are often fused together creating a situation where there is no second contact wedge. Or, in other cases there may be multiple contact wedges oriented in a manner that does not facilitate proper segmentation. A good example of this is shown in Figure 4.9A. This is a particularly clean example in which strong grain boundaries are connected to the contact wedges. For the time being, the grain boundaries are ignored and focus is directed towards the contact wedges. The vertex of the contact wedge is considered to be a corner. The corners found at locations 2 and 3 appear to be good candidates to be joined, however, if corner joining is based solely on proximity measures then corners 1 and 2 will be joined. The correct interpretation is clear from the grain boundaries. It is not sufficient to simply join the closest corners; splitting direction is required to give an indication of how to proceed with the segmentation. This can be determined based on the geometry of the contact wedge by

first determining the line that bisects the wedge and then providing a tolerance on the direction of this line, for example, +-5 degrees. This creates a search wedge, as shown in Figure 4.9B and depending on the radius and angle used to define the wedge, corner 2 will either be inside or outside of this search area. Using the same search criteria for corner 1, results in no suitable matches and the grains will remain un-segmented.

It is clear that in the case shown in Figure 4.9 the grain boundaries provide valuable additional information for the segmentation process. If the grain boundaries are included in the contact wedge definition, the corner location shifts from the vertex of the wedge to the endpoint of the grain boundary. Due to the complex grain boundary interactions found in some of these images, it is believed that the most important information regarding splitting direction is contained the final few pixels of the boundary line. Therefore, extrapolating the grain boundaries based on the last 2 or 3 pixels provides a good estimate of the missing grain boundary. At the same time, the complexity associated with the grain boundaries also implies that there is a limit to the boundary extension.

There are difficulties associated with including the grain boundaries since their presence is dependent on the filtering and thresholding processes that are used in the primary segmentation step. The near continuous boundaries from the color image in Figure 4.10A appear broken in Figure 4.10B after the filtering, thresholding and blob removal operations.

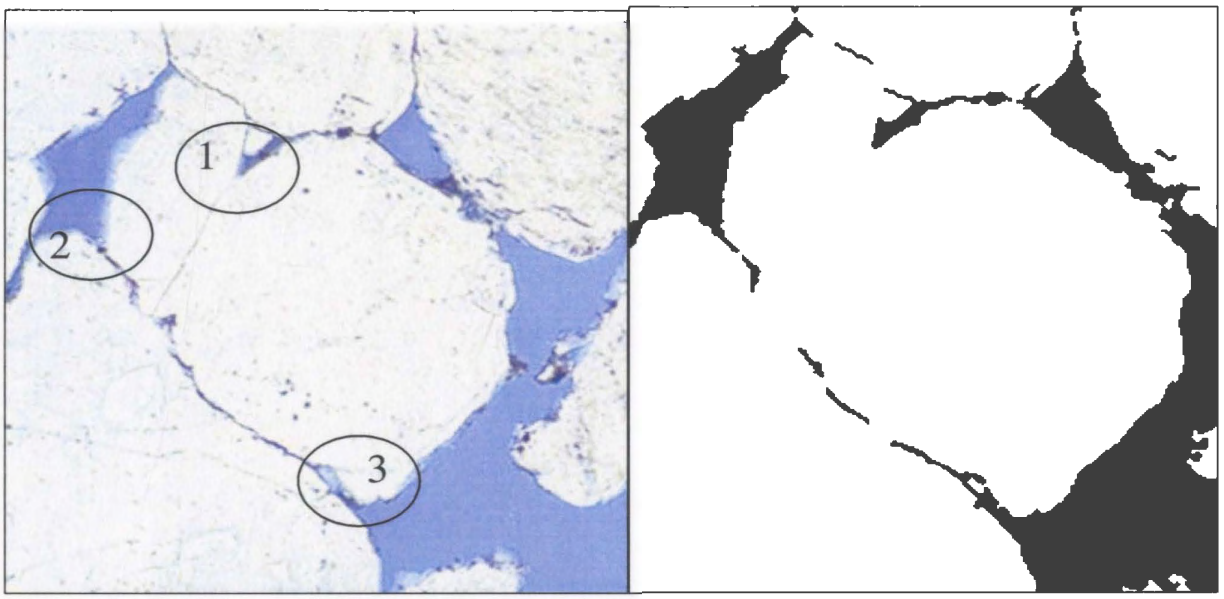

Figure 4.10: Effects of primary segmentation. A) Original image. B) Binary image.

This shows the importance of carefully eliminating blobs from within the grains since these partial boundary segments would have likely been removed if not for the compactness criteria used in the primary segmentation routine. To accommodate these new broken segments, the definition of a corner is changed so that the endpoints of these segments are also considered corners. So, the formal definition of a corner for this work includes contact wedges, endpoints of boundaries connected to contact wedges, and endpoints of broken boundaries. Corner detection schemes were investigated to determine their suitability for finding corners, as defined by the above criteria.

4.2.1 Corner Detection

Corners can be detected using binary images or they can be detected directly from the grayscale images. The digital cutting method (DCM) [van den Berg et al., 2002] is an example of the binary approach and it uses a boundary tracking routine to find suitable contact points that are considered to be corners. The SUSAN detector [Smith and Brady,

1997] is an example of a routine based on the grayscale image. It searches for patterns in the intensity values that indicate corner locations. The result of applying the SUSAN routine is shown in Figure 4.11 where the black dots indicate corner locations.

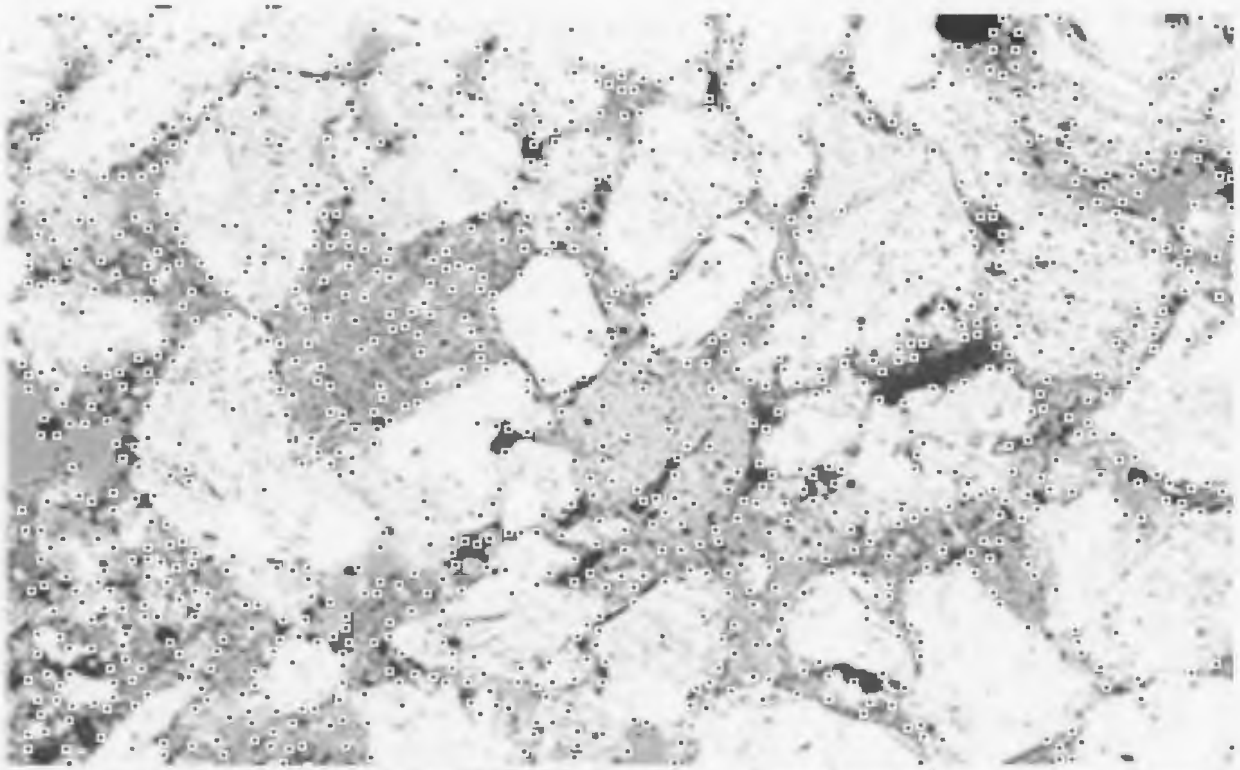

Figure 4.11: Application of the SUSAN algorithm where detected corners are shown by small black dots surrounded by white.

A large number of corners points are detected and only a small fraction of those are considered to be important in the grain splitting process. Such a larger number of corners can be expected for these images due to an abundance of high texture areas. It is possible to tune the routine to produce a lower number of corners but the result is still not acceptable. The algorithm can also be applied to binary images as shown in Figure 4.12. Even though the binary image is much simpler, there are still too many corners detected and it is clear that an alternate approach is required to find the corners needed for this application.

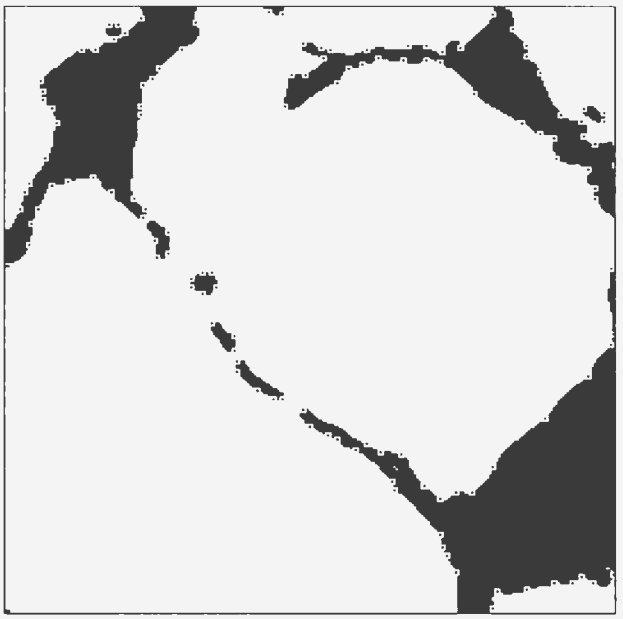

Figure 4.12: Binary image with application of SUSAN detector. Corners shown by black dots surrounded by white.

A new method was developed based on the boundary of the grains in the binary image. The image is dilated 3 times and the difference is taken between this new image and the original image. The result is thinned to skeleton to produce a single pixel wide boundary image. This process is illustrated in Figure 4.13.

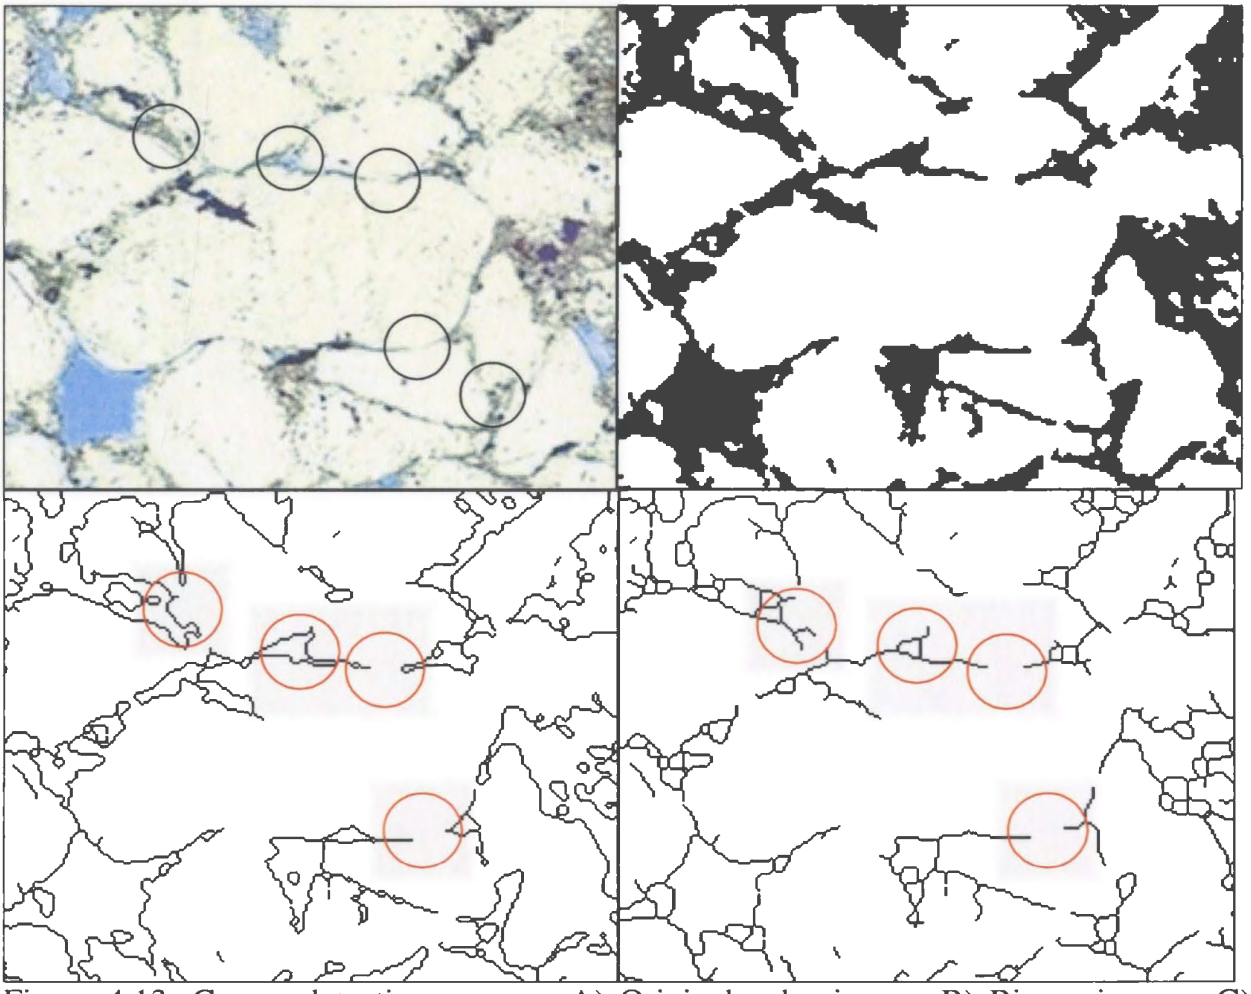

Figure 4.13: Corner detection process. A) Original color image. B) Binary image. C) Skeleton image produced by one dilation. D) Skeleton produced with three dilations.

One dilation operation outlines the grains and produces an image similar to what would be obtained with an edge detector, shown in Figure 4.13C. This is of no real use since a line tracking routine is still needed to find corners. Three dilations will fill most of the gaps between grains, including some contact wedges, and this allows the thinning routine to produce a single pixel wide line, Figure 4.13D, that can be used to indicate corner locations. The focus now shifts to finding and connecting the endpoints. This is accomplished using a combination of line growing and corner connection routines.

4.2.2 Line Growing

Line growing is used to extend endpoints found in the corner detection step. Originally this was achieved using a set of eight 5x5 convolution masks. This is a form of template matching where any pixel that is connected to only one other pixel is selected as a corner and a tangent line is extended based on the orientation of these two pixels. If a line extension segment intersects any other contour in the image, then the line extension stops. If the line extension does not intersect another contour, then the segment is removed using a line shrinking routine. It is expected that this approach will provide quick and accurate segmentation since the line segments are extended from different locations and do not necessarily need to be within the search wedge. So, for in the example used in Figure 4.10, the three corners have the potential to intersect each other as well as any intermediate segments that may be present, as shown in Figure 4.14.

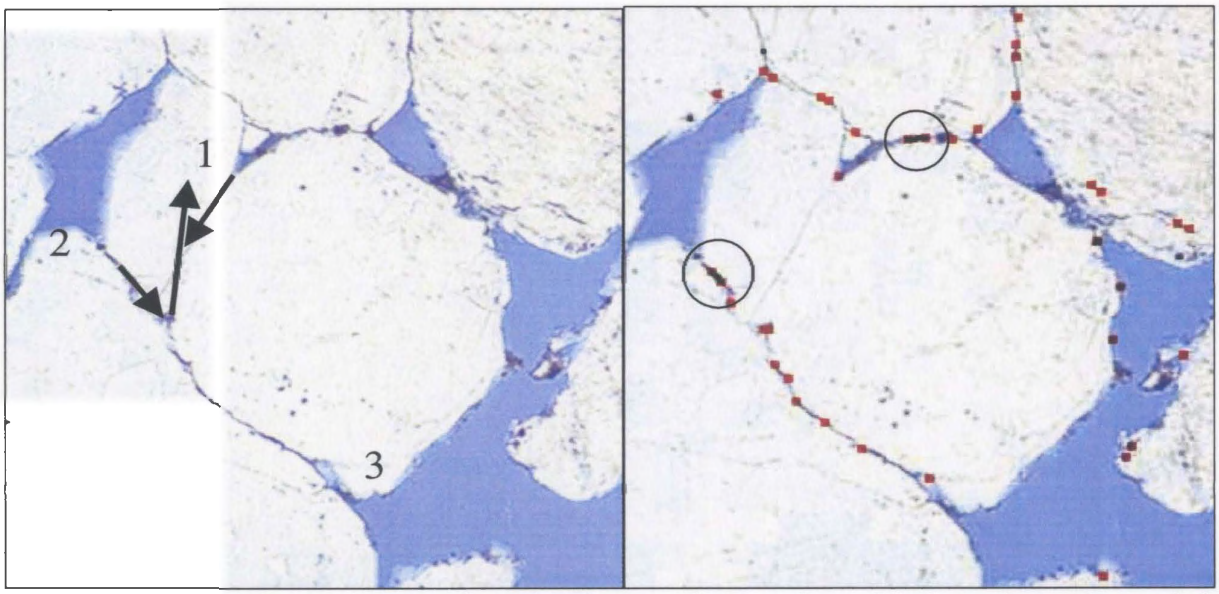

Figure 4.14: Line extension routine. A) Expected operation. B) Actual results, red dots indicated corner locations and two short lines (circled) show successful line extension.

Although this method seems straightforward, Figure 4.14 shows that many corners remain unconnected. This is largely due to complexities in the grain boundaries that cause parallel line extensions that have no possibility of intersecting. Unexpected changes in curvature cause other line extension segments to move in opposite directions.

Using a larger convolution mask partially solves these problems. The slope for the line extension can be based upon 3 pixels rather than two. In order to achieve this, a set of 24, 7x7 convolution masks are used. Using the larger masks reduces the number of corners detected and gives a better approximation of the missing grain boundaries. To adjust for the parallel line segments, a simple corner connection routine is used to join corners that are close to one another. The length of the line extension segments is limited to 5-10 pixels.

4.2.3 Corner Connection

In order to join corners that are close together but not necessarily facing each other, a simple corner connection routine is used. If two corners are detected within a certain distance of each other, they are connected. This routine can be applied as part of the line extension routine or on its own. Good results are observed with this method. The radius of the search area is limited to 5-10 pixels.

4.2.4 Watershed Methods

The corner detection and line growing schemes are limited to contact wedges that are less than 90 degrees. For wedges that are greater than this value, watershed methods are considered. Watershed segmentation was used on this data set by Zhao [2000] and it was

considered to be a poor approach. This is due mainly to the fact that the rough boundaries found in these images cause false watershed. However, if the length of the allowable watershed line is limited, it will not produce as many errors. In fact, small watersheds perform very well for segmenting grains that are in point contact. As with the length of the line extension segments, watershed lines are limited to 5-10 pixels.

4.3 Grain Reconstruction

Partial grains become more abundant as the segmentation progresses, due mainly to the use of boundaries and line segments with lower contrast. These weaker boundaries correspond to quartz overgrowths, boundaries within polycrystalline grains and rock fragments, cleavage planes in feldspars, and in some cases they are scratches left by the sample preparation. The reason for including boundaries found in the lower intensity range is that many of these boundaries are vital to the segmentation process. Considerable effort was expended in an attempt to classify these weak boundaries but no suitable means was found, so instead of attempting to classify every weak boundary, a process was developed to reconstruct over-segmented grains based on shape. Human operators can mentally reconstruct these over-segmented grains and are then able to outline the overall shape of the grain. The goal of the blob reconstruction routine is to mimic this mental reconstruction. Once the segmentation routines have finished, the silt-sized grains and dirt material are added together and analyzed using a combination of morphological operations and blob analysis. If any new blobs are formed they are checked using the grain criteria. If these criteria are met, the new blob will be classified as a grain and added to the removed grain buffer.

Chapter 5

Integration of Methods

The primary goal of this chapter is to integrate the primary and secondary segmentation routines in a manner that optimizes the segmentation. While developing these routines, it became apparent that some methods perform better on some types of grains than they do on others. Each method also requires that local parameters be set in order to optimize the results. One of the most important tuning parameters is the threshold for the initial binarization. The threshold is considered to be a global parameter since changing its value can affect the outcome of every method used. The number of applications of the median filter is another important global parameter and its value will affect all routines. The final factor to consider in the optimization is the order in which the routines are applied.

The secondary goal of this chapter is the development of an automated method for evaluating the performance of the segmentation routine. Originally, the accuracy of each method was determined manually and segmented grains were labeled as either correctly segmented, under segmented, or over segmented. The subjectivity involved in this method was removed by developing an automated accuracy determination routine. First, all routines are tested on a single image with manual inspection; then the routines are combined and tested on 4 images with manual and automatic inspection to ensure that the automatic analysis is performing correctly.

5.1 Individual Routine Performance

Each individual segmentation routine is tested separately in a systematic manner to investigate the affect of each routine and the parameters needed to optimize the segmentation process. This is a time consuming process so only one image was selected for this treatment.

5.1.1 Threshold Selection

The global parameters were the first to be addressed since they affect all the other routines. Based on the discussion of thresholds in Chapter 4, it stands to reason that more than one threshold can be selected to provide an estimate of the grain fabric. Alternate threshold values are obtained by adjusting the high and low correction factors. For the preliminary work the high value was set to 30 and the low value was set at 20. Increasing the correction factors slightly to 35 and 25 resulted in an increase in the number of correctly identified grains. Reducing the correction factors to 20 and 10 also produces an increase in the number of correctly identified grains. In general, increasing the correction factors will give lower thresholds and therefore more pixels are included in the grain fabric. These additional pixels are typically associated with weak grain boundaries and their presence means fewer grains are separated. The opposite is true when lowering the correction factors, the threshold is higher and more of the weak boundary pixels are removed resulting in more partial grains being produced. These results suggest that thresholds selected between these extreme values may show better results and demonstrate a trade-off between under and over segmented grains. However, for this test image the extreme correction factors give better results than any intermediate values.

Although the number of grains being considered is not large enough to make any solid conclusions about this trend, the other segmentation routines, which use the extreme threshold values, consistently show better results than those based on intermediate threshold values. The higher number of pixels being classified as grain material is demonstrated by the fact that for the 35-25 case the grain fabric composes nearly 60% of the image while for the 20-10 case only 50% of the image is grain fabric. This suggests that approximately 10% of the image is composed of pixels that can be considered either grain pixels or boundary/other pixels. It is unclear which set of correction values would provide better overall all results so two algorithms are developed, a high case which uses the larger correction factors and a low case that uses the smaller ones.

5.1.2 Median Filter Applications

Within each case the number of median filter applications is varied from 0 to 5 and the results are observed. The general effect is that one or more applications are needed for the low case and zero applications are required for the high case. Again this fits well with the threshold observations since smoothing the image with the median filter eliminates some of the weak edges that lead to over-segmentation. The high case benefits from having all original boundaries left intact and unaltered by filtering. This analysis facilitates the formation of a high and low base case by which all other routine performance can be compared. The base high case uses the larger correction factors and zero applications of the median filter. The base low case uses the lower correction factors and one application of the median filter.

5.1.3 Other Global Parameters

These cases use three iterations of closing in the blob analysis routine used to detect grains. It was found that three iterations significantly improved the results. Although it was suggested in Chapter 4 as a possible way of assisting in the segmentation, no opening operations were performed on the base case. The opening only improved the high case slightly and it reduced the percentage of correct grains for the low case by almost 10%. This reduction is in the form of over-segmentation.

5.1.4 Individual Results of the Segmentation Routines

After defining a set of base cases the remaining primary and secondary segmentation routines can be investigated. The best results for the Sobel routine, line growing and corner connection routine, and the watershed routine are demonstrated in the following tables. For all of these cases, the result of the segmentation routine is compared to the complete, manually segmented image, i.e. no consideration is given to the different grain types. Focus is placed on methods that provide correct grain percentages in the 75% and greater range. The results of the Sobel segmentation routine are shown in Figure 5.1. The number in the legend refers to the threshold value used to binarize the Sobel intensity image; in this case two thresholds, 60 and 100, were selected. The Sobel detector is applied to a filtered version of the red channel image. The number of filter applications is independent of the number of applications used to produce the binary grain image.

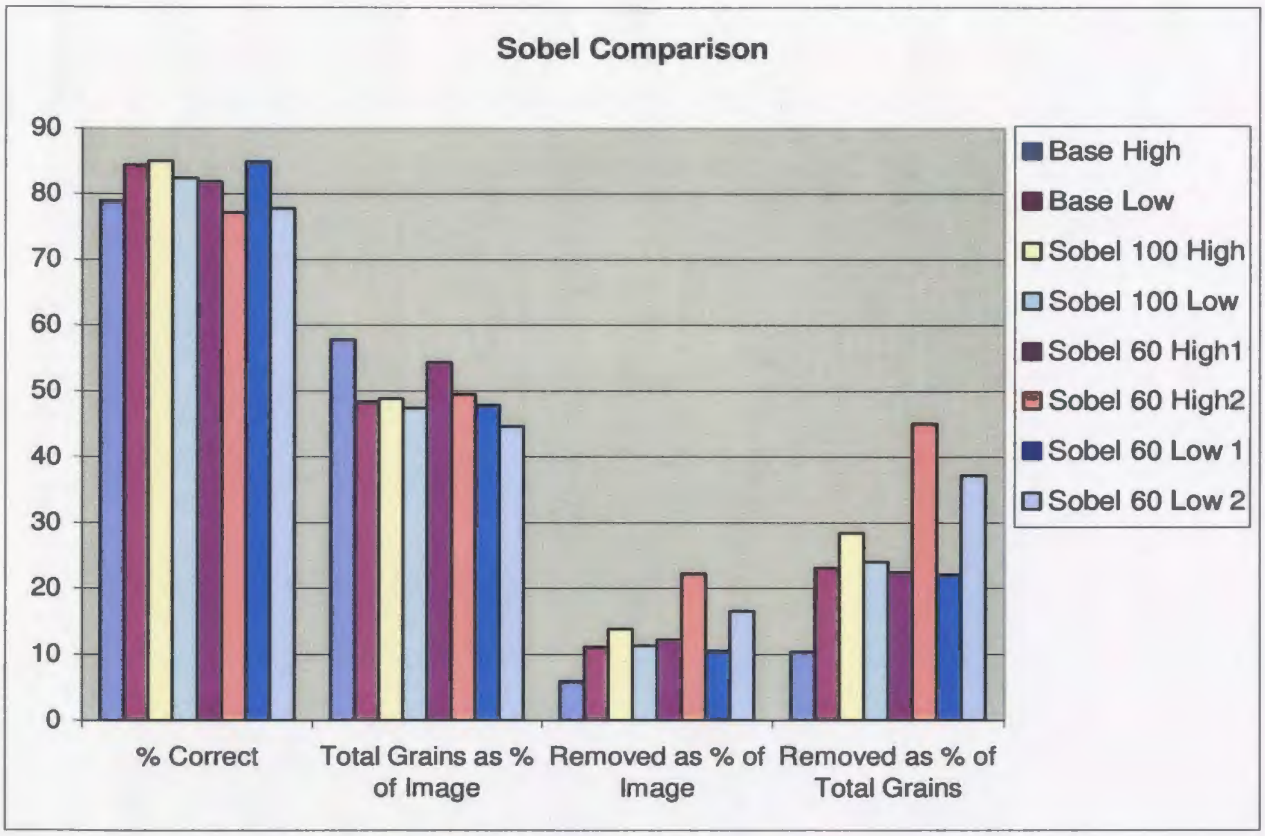

Figure 5.1: Best results for the Sobel routine compared to the base cases.

Pixels above the threshold value are kept and therefore a higher value eliminates weaker boundaries. So for the Sobel 100 case the threshold reduces the need for filtering and peak accuracy is obtained with no filtering in the high case and one application in the low case. The Sobel 60 case requires additional filtering to improve the number of correctly identified grains. The first test with Sobel 60 uses five applications of the median filter and has more correct grains as compared to the second test, which uses just one filter application. For both the high and low case the second test shows a lower percentage of correct grains but a significantly greater number of grains are segmented. Once again the incorrectly identified grains are partial grains. Another important measurement is the total grain fabric present, which is a measure taken after the silt sized particles have been removed from the image. So, for routines that produce lower initial grain fabrics the

amount of segmented silt-sized material is greater. Taking a second look at the Sobel 100 high case shows that it does a good job with the segmentation according to the percentage of correct grains and amount of grain material detected. However, the initial amount of grain fabric is nearly 10% less than the base case, all of which is classified as silt-sized grains or smaller. This can be attributed to the lack of filtering which is an important consideration for selecting segmentation routine parameters.

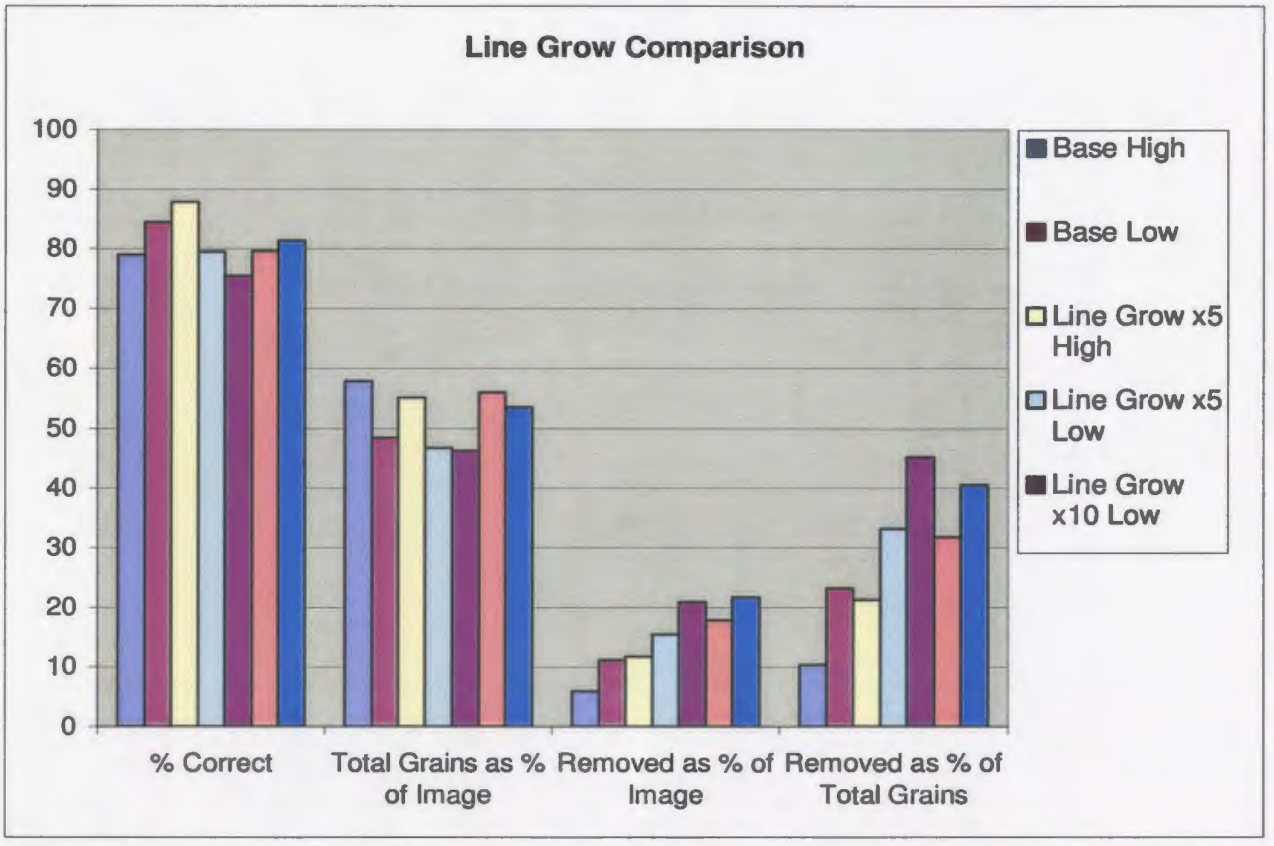

Figure 5.2: Best results from the line growing and corner connection routine compared to the base cases.

The results for the line growing and corner connection routine are shown in Figure 5.2. At this point, the behavior of the line growing and corner connection routines can be predicted based on the above discussion. Specifically, median filtering will affect the number of corners found and therefore the number of line extension locations. The

combination of thresholding and median filtering will affect how many iterations of the line growing routine are needed to split grains. In terms of accuracy, 5 iterations of the line growing routine shows good results for both the high and low case. While this is a good start, the high case removes only 21% of the grains and the low method removes only 33%. Increasing the number of iterations to 10 results in a higher percentage of grains being segmented but at the cost of reduced accuracy.

Adding a corner connection routine with a 5 pixel radius to the line growing gives good results for the high case with 32% of the grains removed at 80% accuracy. However, for the low case the accuracy is the same as the line grow x10 case but it removes only 41% of the grains, down from 45%. The difference between the two methods is the number of filtering applications used; for the line growing case 5 applications are used and for the line growing and corner connection case only 1 application is needed. This suggests that the extra filtering changes the image in such a way that the accuracy is improved for the line growing method. This could be in the form of altering the boundary curvature or just eliminating erroneous line growing locations. Extending the corner connection radius to 10 pixels, once again, increases the number of grains segmented by the high case from 32% to 41% with an accuracy of 81%. Also, by reducing the amount of filtering from 1 application to zero the number of segmented grains is increased to 55% with only a 6% reduction in accuracy. This is a good result but further examination of the percentage of initial grain fabric suggests that the median filtering is needed. Removing it results in a 6% decrease in the initial grain fabric, and hence possible over segmentation. Increasing the corner connection radius has no real affect on the low case and the percentage of segmented grains is 44% with

74% accuracy. The results for the low case seem to suggest that a limit exists for the number of grains that can be removed using the contact wedge approach. In other words, only a fraction of the grains are connected in a manner that is conducive to segmentation by the line growing based routines. This is confirmed by visual inspection of this particular image.

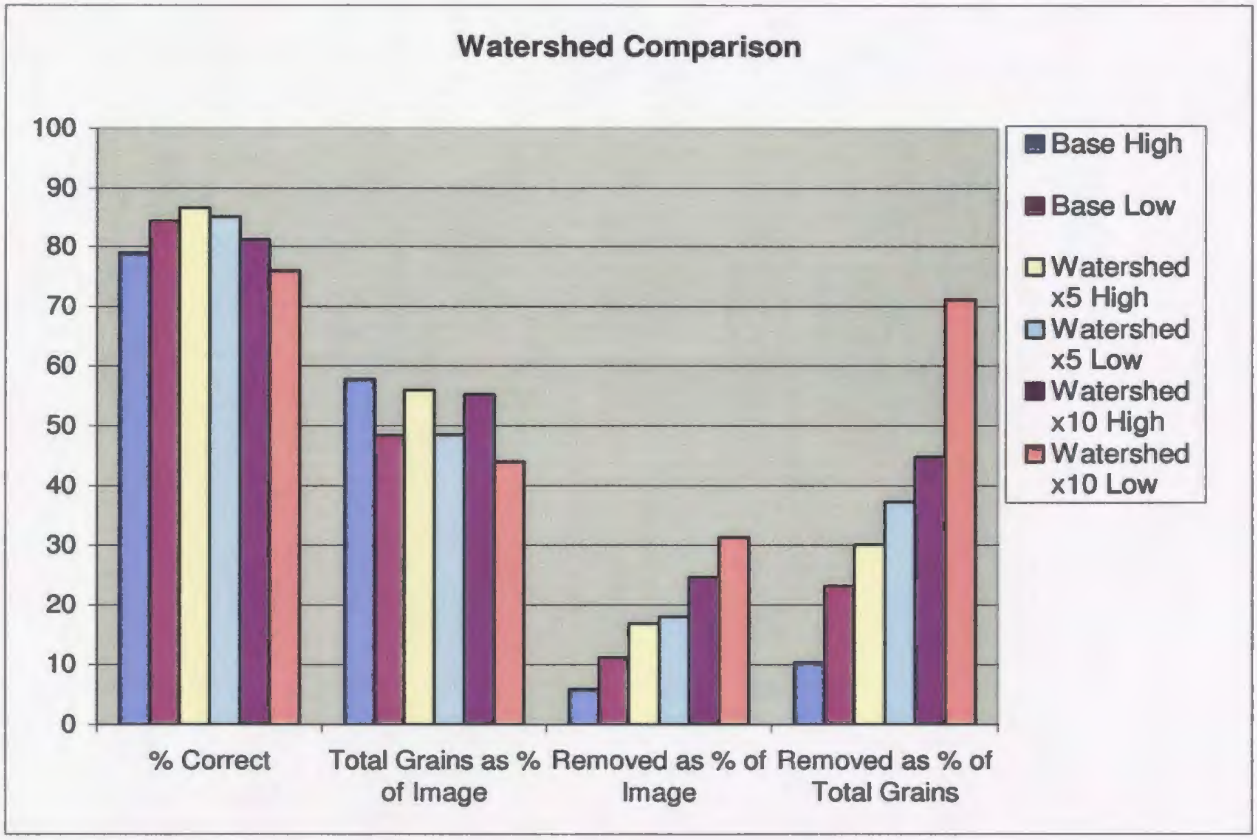

Figure 5.3: Best results for the watershed routine compared to base cases

The result of applying the watershed routine is shown in Figure 5.3. The cases tested select watershed lines that are less than 5 or 10 pixels in length. The watershed routine is less sensitive to variation in thresholds and filtering and hence, less sensitive to noise. Also, they show a significant improvement in segmentation as compared to the base cases. The watershed x10 low case segments over 71% of the grain fabric with an

accuracy of 76%. The lower correct grain percentage indicates some over segmentation, also indicated by the low initial grain fabric percentage. The 10 pixel watersheds have a lesser affect on the high case since there is more grain material present.

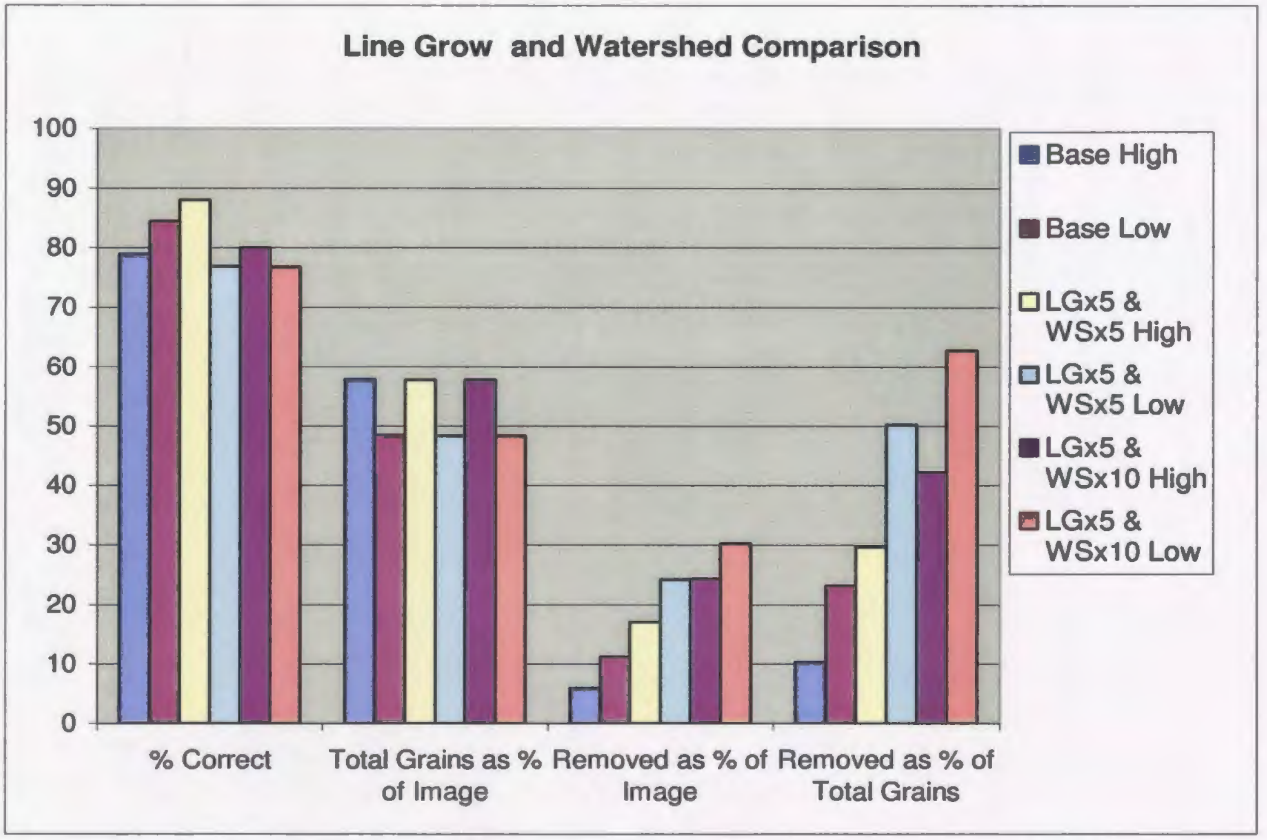

Figure 5.4: Best results for the line growing and watershed routine compared to the base cases.

Finally, results are shown in Figure 5.4 for the line growing and watershed method. For the high cases, the results are nearly identical to those obtained with the line growing and corner connection routine. The exception being, this routine does not require any filtering while the line growing and corner connection requires one application. This agrees with the observations made previously about the watersheds relative insensitivity to noise. For the low cases, the accuracy is slightly lower than the high case but slightly better than those obtained for the low cases of the line growing and corner connection

methods. Also, there are significantly more grains segmented with no reduction in initial grain fabric, suggesting few silt-sized grains are created as a result of over-segmentation.

5.2 Combining Methods

Now that there is some insight into the individual behavior of each of the methods, focus shifts to combining routines in a manner that will maximize segmentation. The strategy is straightforward; the most accurate methods are applied first. Also methods with shorter line lengths are applied before those with longer lengths, for example, the watershed x5 routine would be applied before the watershed x10 routine. In theory, the secondary segmentation methods should have some overlap in the types of grains that they segment and therefore only a small improvement is expected from routines that are applied later on in the algorithm. For the Sobel methods, the higher threshold is applied first since it includes stronger edges and therefore causes less over-segmentation. Two cases are developed based on the high and low cases discussed in the previous section. These cases are summarized in Tables 5.1 and 5.2.

High Case - Median x0			
Order	Type	Accuracy	Rank
1	Thresh Mx0	79	Max
2	Watershed x5	87	Max
3	LineGrow x5	88	Max
4	Sobel 100 Mx0	85	Max
5	Sobel 60 Mx5	82	Max
6	Watershed x10	81	Max
7	LineGrow x5, CCx5	78	Compared to 80
8	LineGrow x10	72	Max
9	LineGrow x5, CCx10	75	Compared to 81

Table 5.1: High case definition based on zero median filter applications

Low Case - Median x5			
Order	Type	Accuracy	Rank
1	Thresh Mx5	79	Compared to 84
2	Watershed x5	85	Highest
3	Sobel 100 Mx5	82	Tied for highest
4	Sobel 60 Mx5	85	Highest
5	LineGrow x5	78	Compared to 80
6	Watershed x10	75	Compared to 76
7	LineGrow x10	75	Highest
8	LineGrow x5, CCx5	71	Compared to 75
9	LineGrow x5, CCx10	74	Highest

Table 5.2: Low case definition based on 5 median filter applications

With the exception of the Sobel operators, all methods are required to use the same number of median filter applications since they are based on the same binarized image. The rank column in Tables 5.1 and 5.2 indicates where the selected method ranks compared to the other methods with different amounts of filtering. In most cases, only a small amount of accuracy is given up in order to accommodate for using the same base binary image for all routines. For the high case, the Sobel 100 routine does not use any filtering while the Sobel 60 routines uses 5 applications. Also, the combined line growing and watershed method is not included at this point in the interest of simplicity. In both cases, the accuracy decreases with increasing distance between segmentation locations. The secondary watershed x5 case is included before the primary Sobel case since it takes the place of the opening operation that was proposed in Chapter 4.

The High and Low algorithms are applied to a total of four test images. The segmented grains are then compared with their manually segmented counterparts. The results are compiled manually to gain insight into the interaction of the methods and to provide a benchmark for testing with the automated analysis algorithm developed in the next section. After applying the low case to one of the test image many grains were left

un-segmented despite appearing as if they should be separated. This prompted the reduction of filter in the low case from 5 applications to 1 application. The overall accuracy remains the same but the number of segmented grains rises from 30% to 101%. This result is shown in Figure 5.5 along with the results from the high case.

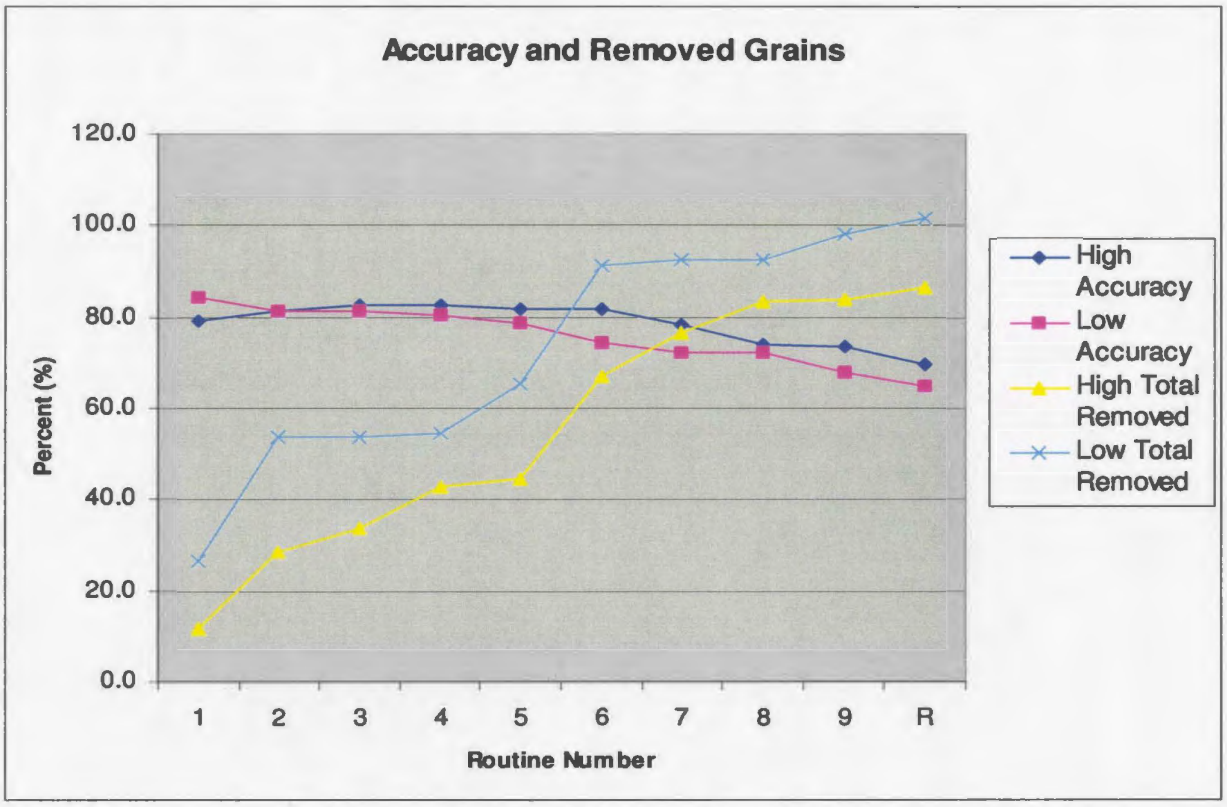

Figure 5.5: Accuracy and Removed grains for one of the test images.

The accuracy for both the high and low case is very similar and decreases as the algorithm progresses. The cause for the abnormal percentage of removed grains is based on the fact that some grains are reconstructed from silt-sized grains that are not included in the original estimation of the grain fabric area. In addition to this, the closing routine used to classify the blobs tends to slightly increase the area of each blob. Therefore, the area of removed grains is greater than the actual value, however, this is not considered to be significant. The results for the other three images are not as promising in terms of

accuracy and grains segmented. However, for the other images the segmentation algorithm provides better estimates of the average grain area and standard deviation.

The average grain area and standard deviation are computed at each stage of the segmentation and the results for one of the test images are plotted in Figures 5.6 and 5.7 respectively. The same results are included for the manually segmented images.

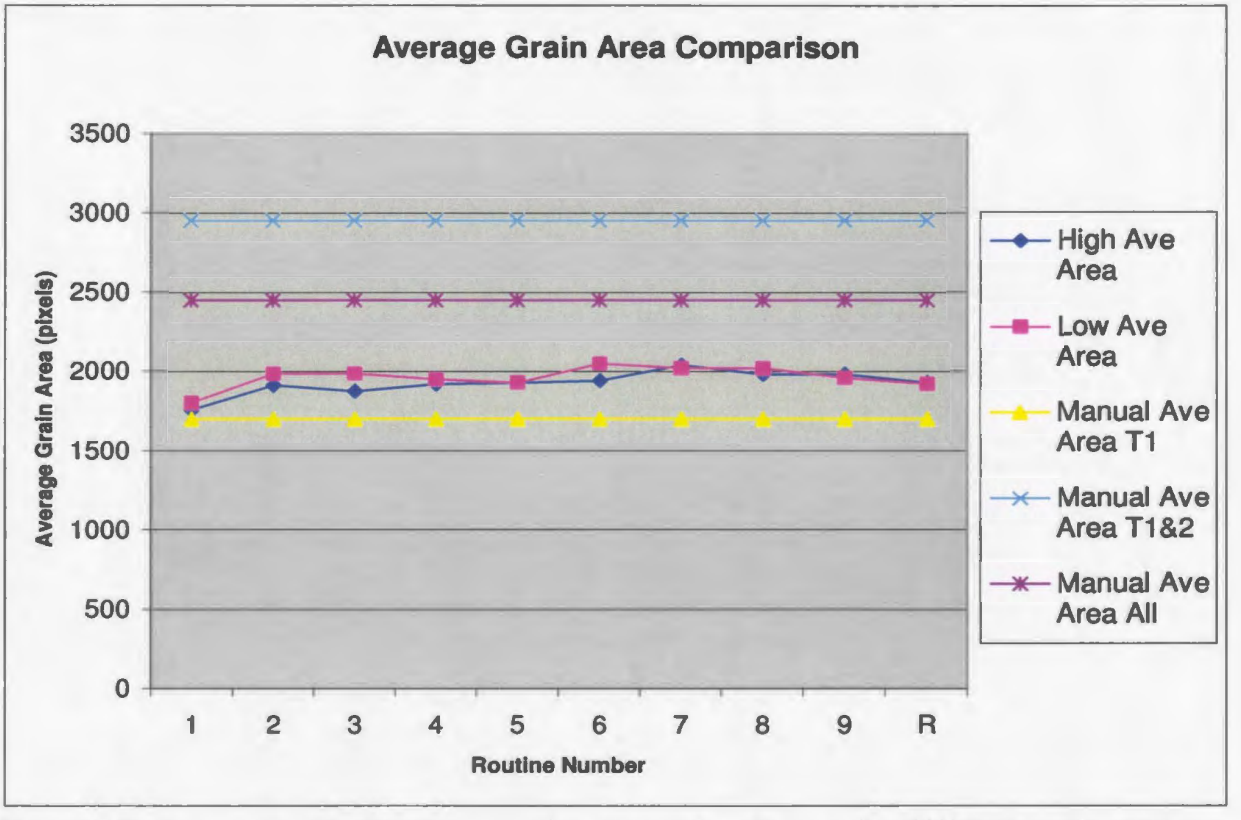

Figure 5.6: Average grain area compared to manual grain area for different grain types.

The average grain area shows considerable variation for each of the different grain type combinations. Ideally, the estimate should be somewhere between the type 1&2 and the 'all types' case. It is interesting that the average grain area for both the high and low case remains very close throughout the algorithm. This suggests that for this image the 10% difference in initial grain fabric area is due to the presence of boundary pixels and not

grain pixels. If 10% of the grain pixels were removed, a lower average grain area could be expected. For other images, the experimental average grain area is much closer to the manual grain areas as demonstrated in Figure 5.7.

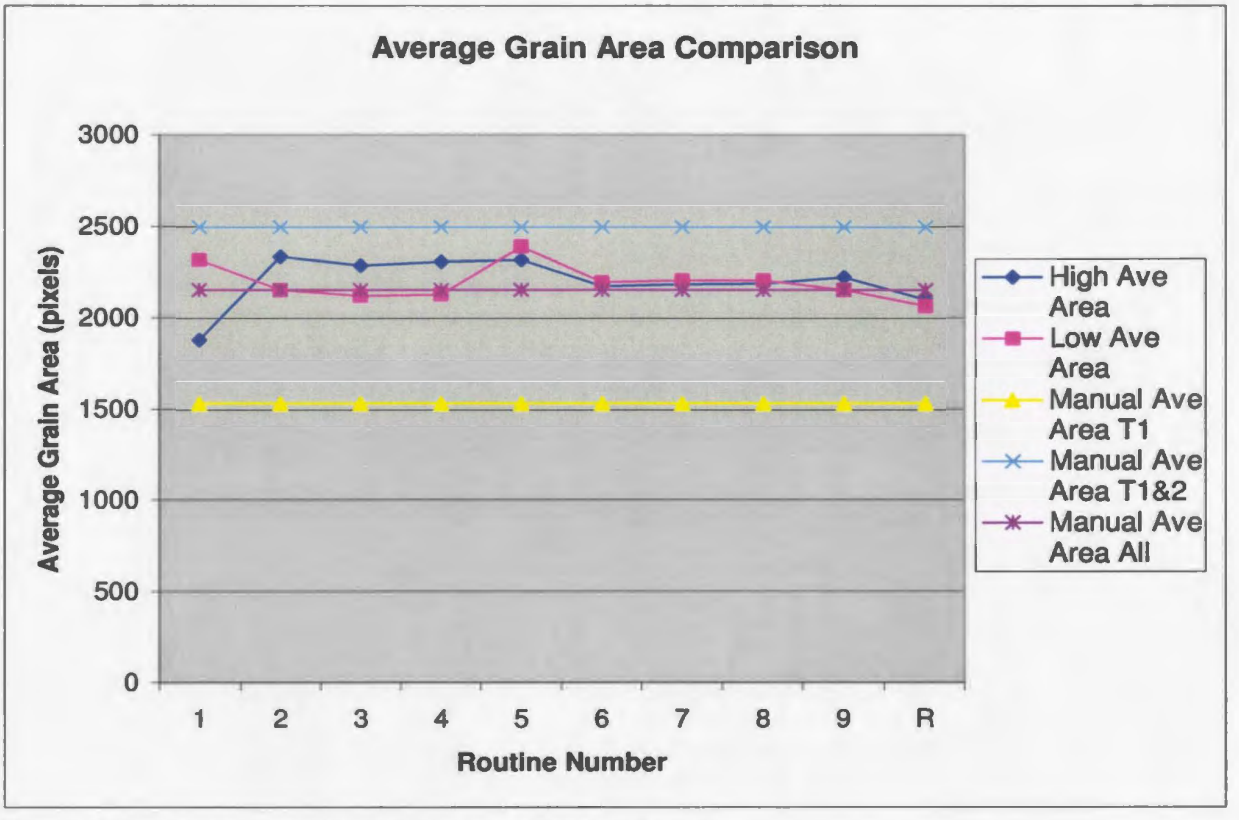

Figure 5.7: Average grain area for another test image.

Again both the high and low cases give roughly the same average grain area throughout the algorithm. Also, the values fall between the type1&2 and 'all types' value, with the final values converging on the 'all types' case. There are a number of very large and very small type 3 grains in this particular image causing the standard deviation to be increased when they are included in the analysis. These type 3 grains are not found using the segmentation algorithm as demonstrated by the standard deviation, shown in Figure 5.8, that is in the range of the type 1 and type 2 grains.

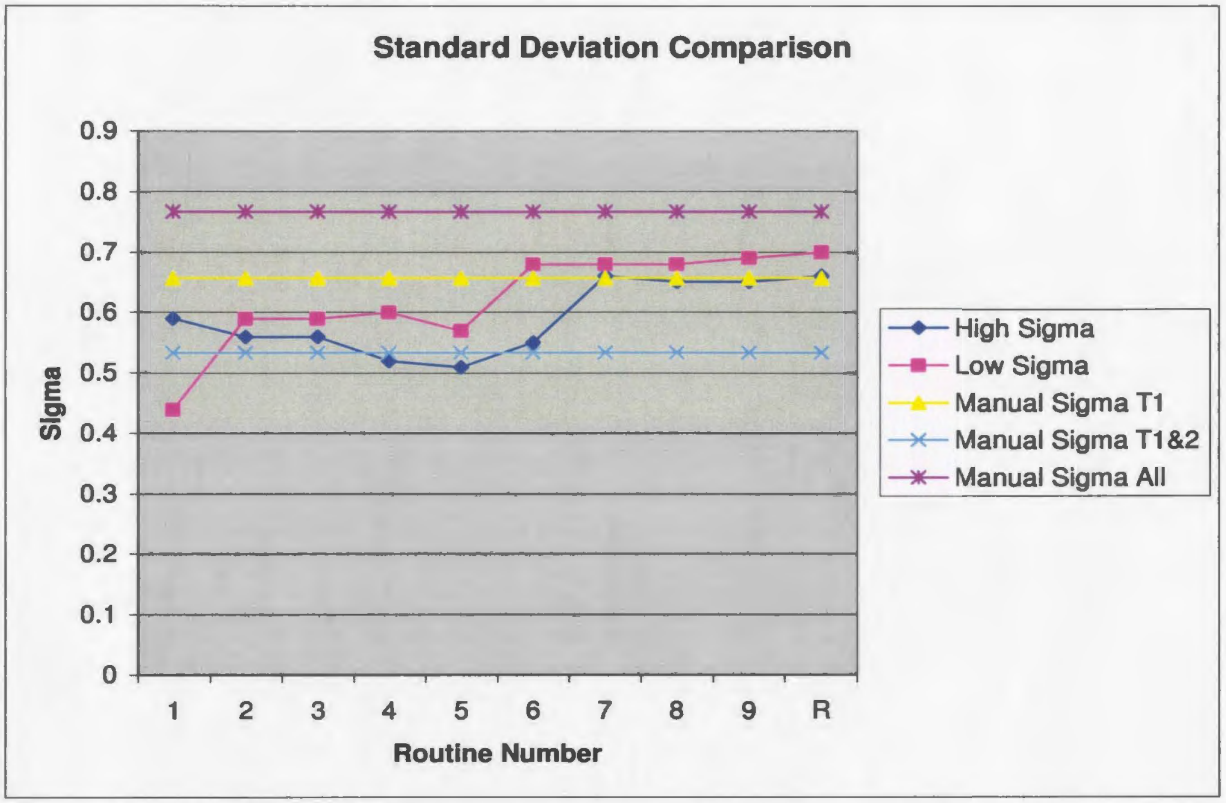

Figure 5.8: Standard deviation compared to manual standard deviation for different grain types.

5.3 Automated Accuracy Routine

Compiling results of the automated segmentation routine is very time consuming and subjective. These drawbacks prompted the development of an automated analysis routine that can compare the results of the segmentation process to the manually segmented images. This method is based on the centers of gravity (COG's) of the automatically and manually segmented grains. The first step is to check for COG's from the manual grains within the boundaries of the potential grains. If two or more manual COG's are found within the potential grain, it is then considered to be under-segmented. If two or more automatic COG's are found within the manually segmented grain they are then considered to be fragments of an over-segmented grain. If there is only one COG, the

areas of the two are then compared and if the area of the automatic grain is within 20% of the manual grain area, it is then considered to be correct; if not, it is considered to be either under or over-segmented. The output of this routine is the original image with the correct, partial, and under-segmented grains displayed on the image, along with the COG's of both automatic and manual grains. This allows for a visual inspection of the performance of the segmentation algorithm. Also, all the statistics discussed in the previous section are output to text files for further analysis. To ensure that this routine is performing correctly, the results are compared to the manual analysis carried out on the four images in the previous section. These results are shown in Figure 5.8.

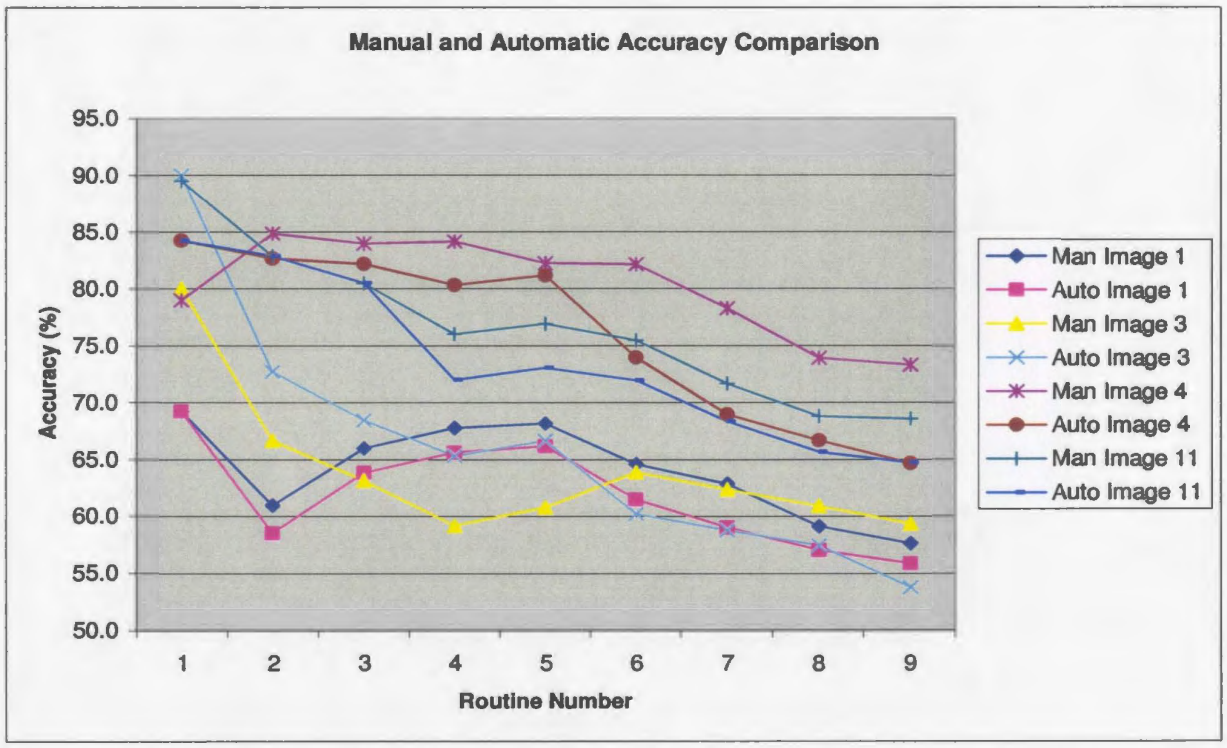

Figure 5.9: Comparison of manual and automatic analysis for the four images examined in the previous section.

Some discrepancies are noted between the automatic and manual analysis specifically for image 4 which is off by nearly 10%. Visual inspection attributes this to subjectivity in

the manual analysis and the performance of the automated analysis is considered to be quite good overall.

Chapter 6

Results and Discussions

The discussion of results is initiated by illustrating specific examples resulting from applying the Low Case Mx5 routine to the first image in the group A set. Next the other cases are applied to the same image and the results are compared. Once this is done, the results of all four cases are shown for all 14 images in group A. Finally, results from Group B are discussed and some comments are given regarding overall program functionality

6.1 Group A Results

The original image is shown in Figure 6.1 while Figure 6.2 displays the result of the segmentation process and Figure 6.3 displays the output from the comparison to the manually segmented image, shown in Figure 6.4. The low case with five applications of the median filter (Low Case Mx5) is used since it uses the highest threshold values and the most filtering of all the cases considered. The results of this case should provide the highest number of correctly identified grains but with the lowest amount of segmented material. Table 6.1 provides a legend to aid in the interpretation of the output images.

Segmentation Ouput		Accuracy Output	
Color	Meaning	Symbol	Meaning
Blue	Potential Grains	C	Correct grains
Light Blue	Silt sized grain material	P	Partial grains
Black	Pore space	US	Under-segmented grains
Pink	high texture areas	Red dot	Location of manual center of gravity
Yellow	Opaque Material	Green dot	Location of automated center of gravity
Green	Reconstructed Material		
Red	Very small misc components		

Table 6.1: Legend for interpretation of output images.

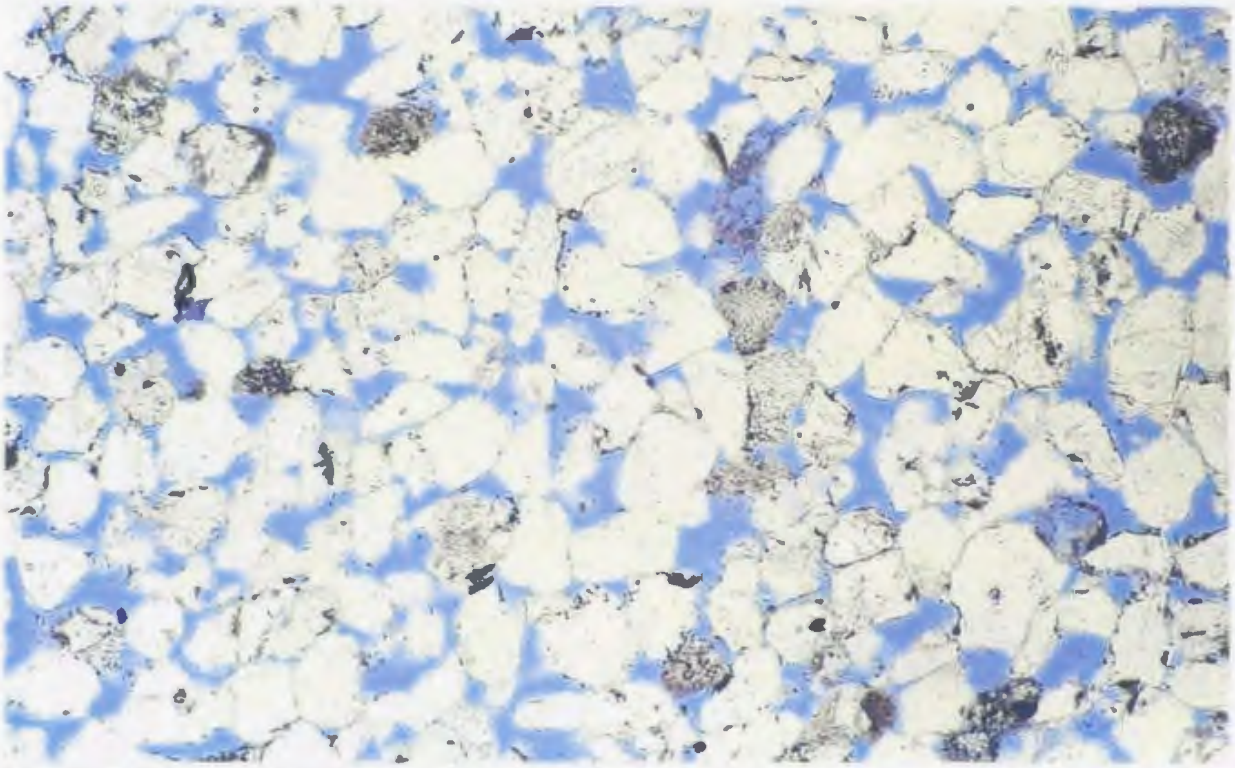

Figure 6.1: Original Image

The algorithms performance is considered by examining the correct, partial, and under-segmented, individually. In addition, an explanation of un-segmented material is given for some cases.

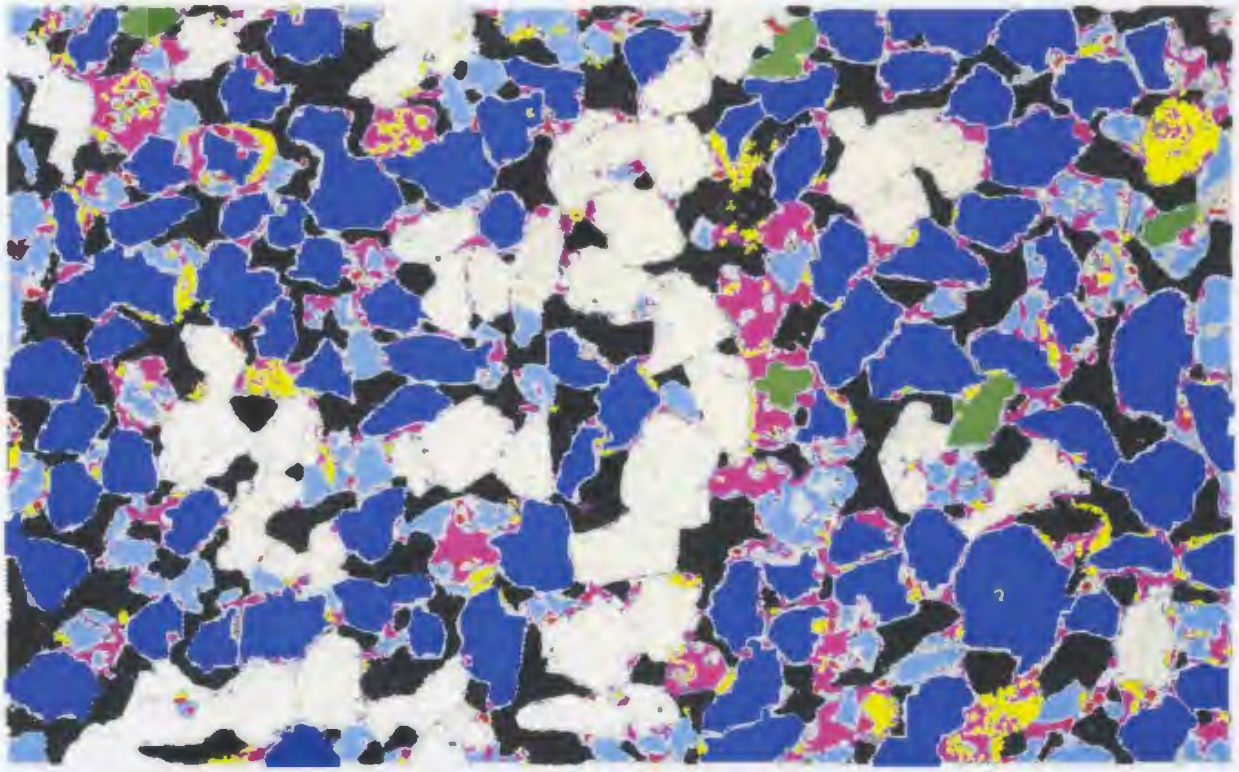

Figure 6.2: Output image from segmentation algorithm - Low Case Mx5.

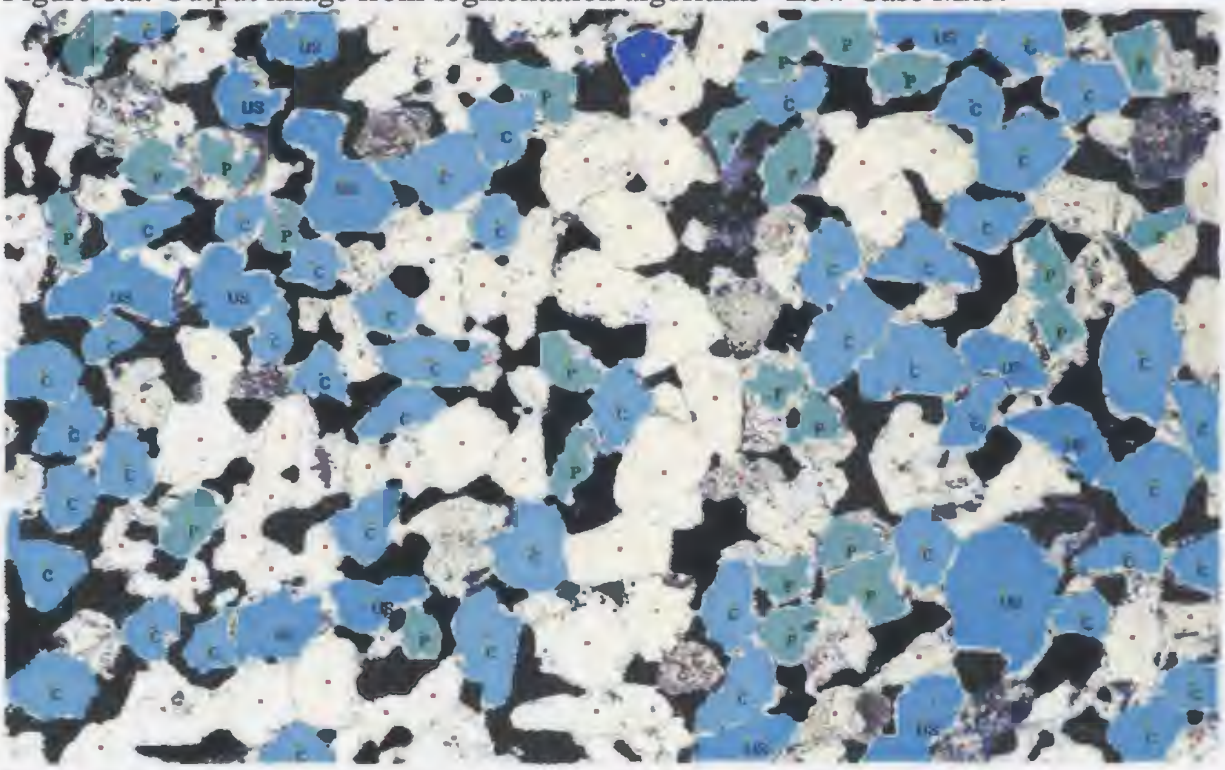

Figure 6.3: Output from manual image comparison – Low Case Mx5

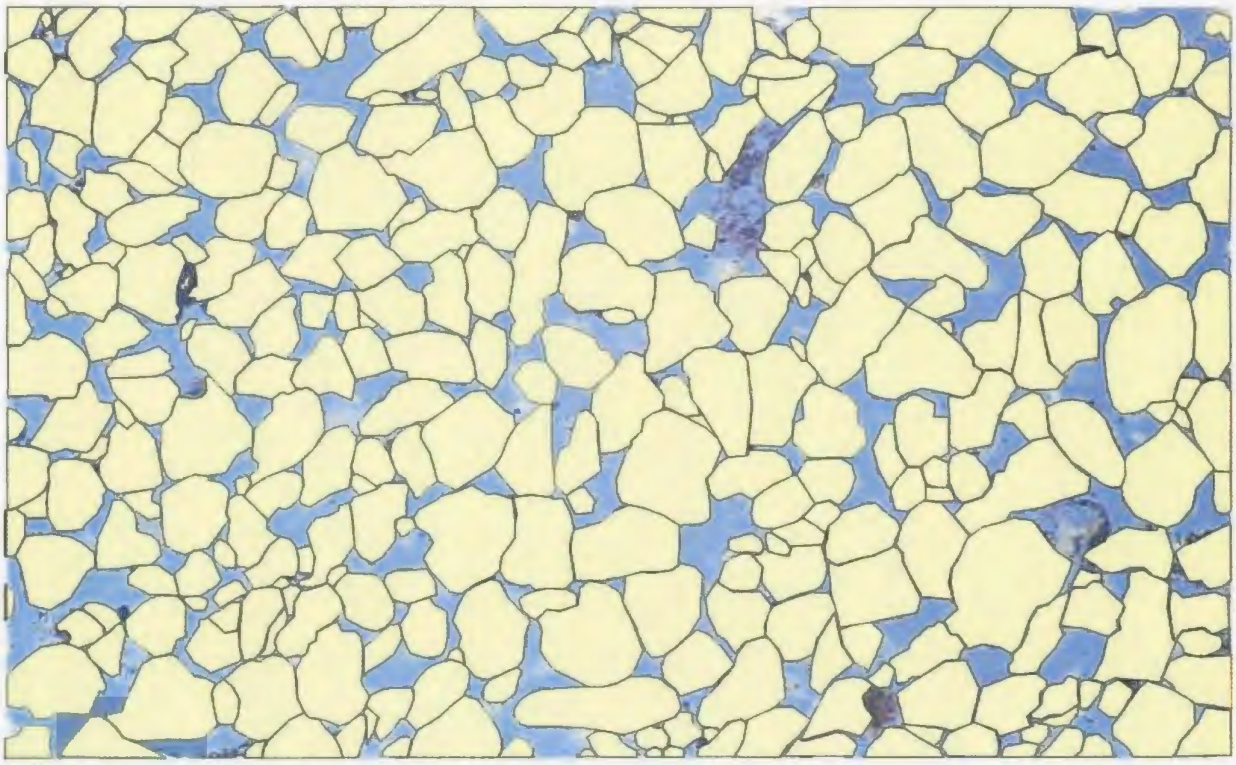

Figure 6.4: Manually Segmented image

6.1.1 Correct Grains

Initial examination suggests the correct grains seem to agree well with the manually segmented images, however, on closer inspection there are cases where grains that are identified as being correct actually overlap with other grain regions. An example of this is shown in Figure 6.5 where A is the original image, B shows the manual outline and C is the result of the automated segmentation algorithm.

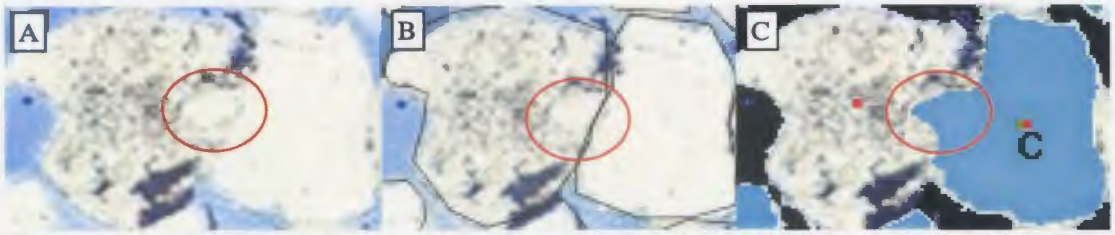

Figure 6.5: Overlap of correct grains. A) Original. B) Manual outline. C) Automatic result.

The area of overlap (indicated by red oval) should not be included as part of the grain. A more extreme example is given in Figure 6.6.

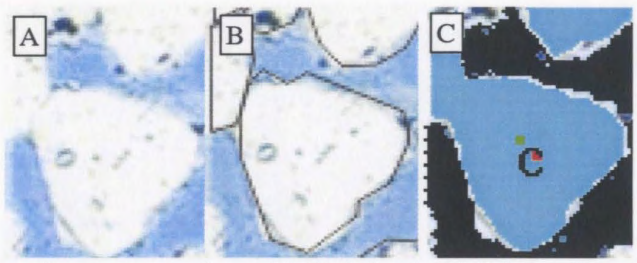

Figure 6.6: Overlap of correct grains. A) Original. B) Manual outline. C) Automatic result.

These types of errors are unavoidable since some allowance must be made for the fact that automatically segmented grain will seldom be an exact match with the area of the manually segmented grain. Also, the grain identification criteria are not perfect and while it may be possible to reduce errors like those shown in Figure 6.6 by introducing additional criteria, the errors shown in Figure 6.5 are difficult to identify. The opposite outcome is illustrated in the next section.

6.1.2 Partial Grains

Partial grains are typically associated with high texture regions of the image or polycrystalline grains. In some cases, as displayed in Figure 6.7, the partial grains are only slightly smaller than the manual grain and the difference in size is the result of small pieces of grain being 'chipped away' from the perimeter of the grain.

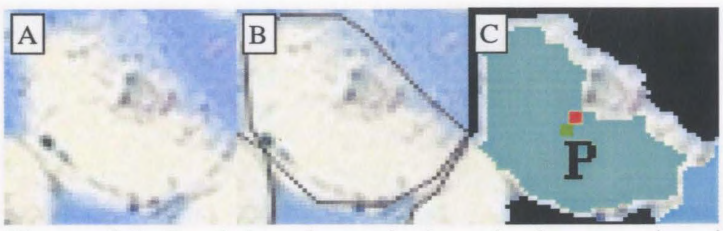

Figure 6.7: Partial grain - missing pixels around majority of perimeter.

Cases like this are dependent upon the operators choose of manual boundary, which is often arbitrarily selected, especially in cases where shelving effects are pronounced. In other cases, as displayed in Figure 6.8, the partial grain is only a fraction of the complete grain.

Figure 6.8: Partial grain – significantly reduced area.

The high threshold used for this case contributes to the effects seen in the above images since there is a greater number of low intensity pixels eliminated from the image. These types of grains are classified as type 3.

6.1.3 Under-segmented Grains

Two main types of under-segmented grains are found; those that are composed of two or more smaller distinct grains, as shown in Figure 6.9, and those that are composed of grains that simply blend together with other grain regions, as displayed in Figure 6.10.

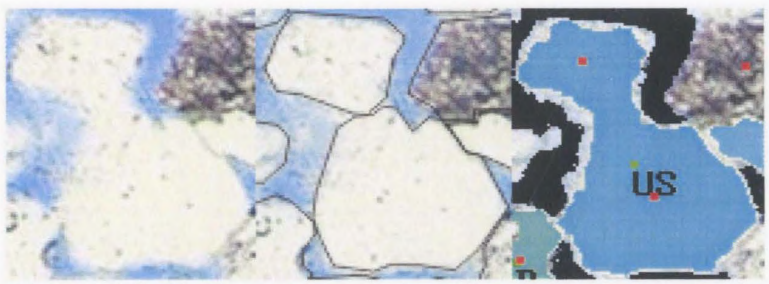

Figure 6.9: Under-segmented grain – two distinct grains

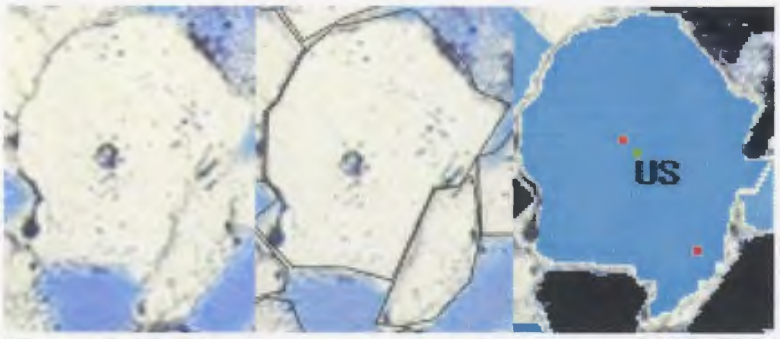

Figure 6.10: Under-segmented grain – one grain blends into the other

These examples are similar to those shown in Figure 6.5 and 6.6 with the difference being that they do not meet the grain criteria. As discussed previously, errors such as the one in Figure 6.9 can be reduced while those shown in Figure 6.10, considered to be type 3, are unavoidable.

6.1.4 Un-segmented Material

There are two main reasons for un-segmented material in this particular image. The first involves the lack of acute contact wedges as shown in Figure 6.11A. Many grain regions have contact wedges that are 90 degrees or greater and the corner detector used in this work is simply not designed to detect those types of corners. The second cause of un-segmented material is the lack of strong grain boundaries, illustrated in Figure 6.11B.

Figure 6.11: Unclassified material – A) Obtuse contact wedges. B) Missing boundary

6.1.5 Reconstructed Grains and Other Material

Five reconstructed grains (green) can be seen in Figure 6.2 and checking these in Figure 6.3 shows that two of these grains are identified as being correct. Also in Figure 6.2, mixtures of pink and yellow material partially identify high texture grains.

6.2 Comparison to Other Cases

The results for the three other cases are shown together with the Low Case Mx5 in Figure 6.12. This graph shows the cumulative percentage of segmented grain material as the algorithm progresses. It is noted that in both the high case and the low case the median filtering causes a slight reduction in the total amount of grain material segmented.

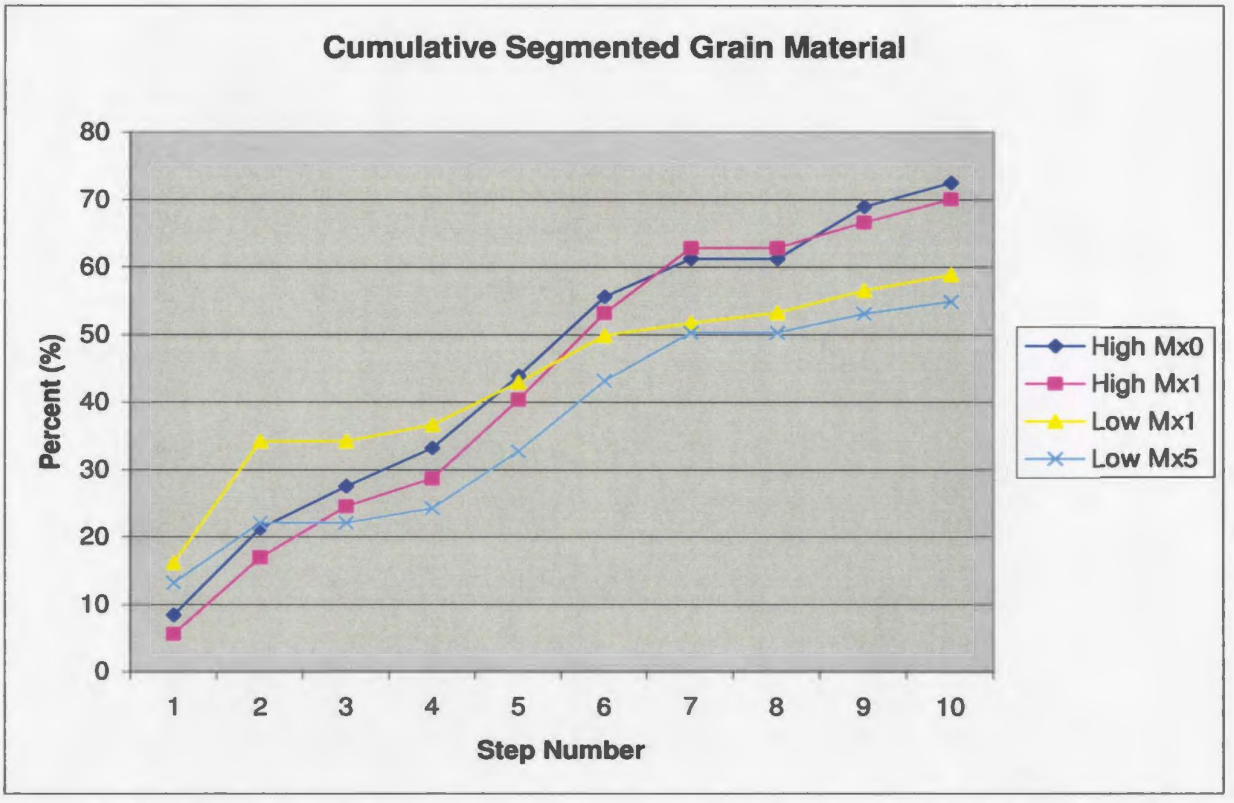

Figure 6.12: Segmented grain material for each of the 4 cases as the algorithm progresses.

The accuracy with which the above grain material is segmented is the focus of Figure 6.13. One extra case, which compares the potential grains to a manual image that includes only type 1&2 grains, is added for comparison. The accuracy for all cases decreases as the algorithm progresses and at the completion all values are roughly the same with the exception of the type 1&2 case. It is theorized that as the algorithm progresses the number of type 3 grains that are improperly segmented increases and, therefore, if these grains are not being considered then the overall accuracy will be better. It is also worth noting that the percentages displayed in Figure 6.13 are based on the numbers of correct grains and if the area of correct grains is considered instead there is roughly a 2% increase in correctly identified material in all 4 original cases. However, in the type 1&2 case there is nearly a 10% increase in correctly identified grain material with the final values being 78%.

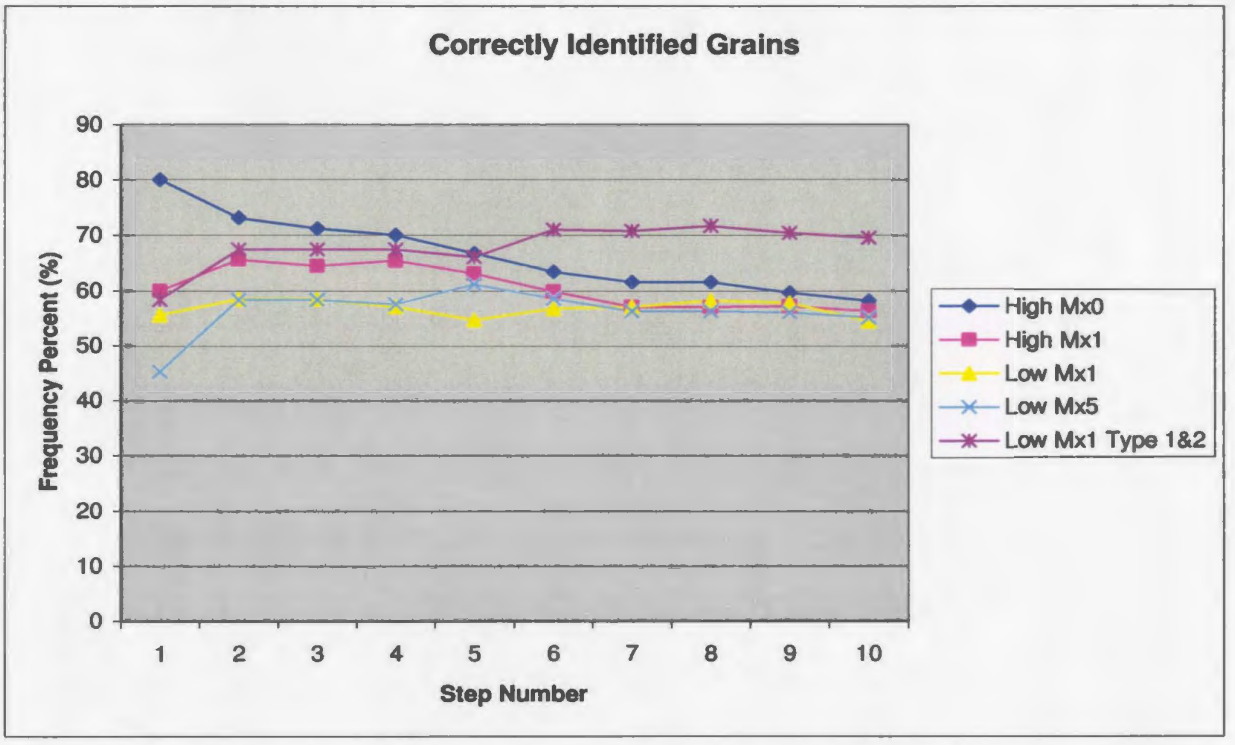

Figure 6.13: Correctly identified grains for each of the cases as the algorithm progresses.

6.3 Comparison to Other Images

The results for the remaining images are compared in this section. The first item considered is the final percentage of segmented grain material for each of the group A images as shown in Figure 6.14. It is important to keep in mind that most images, on average, have only 65 – 70% of the image classified as type 1 and type 2.

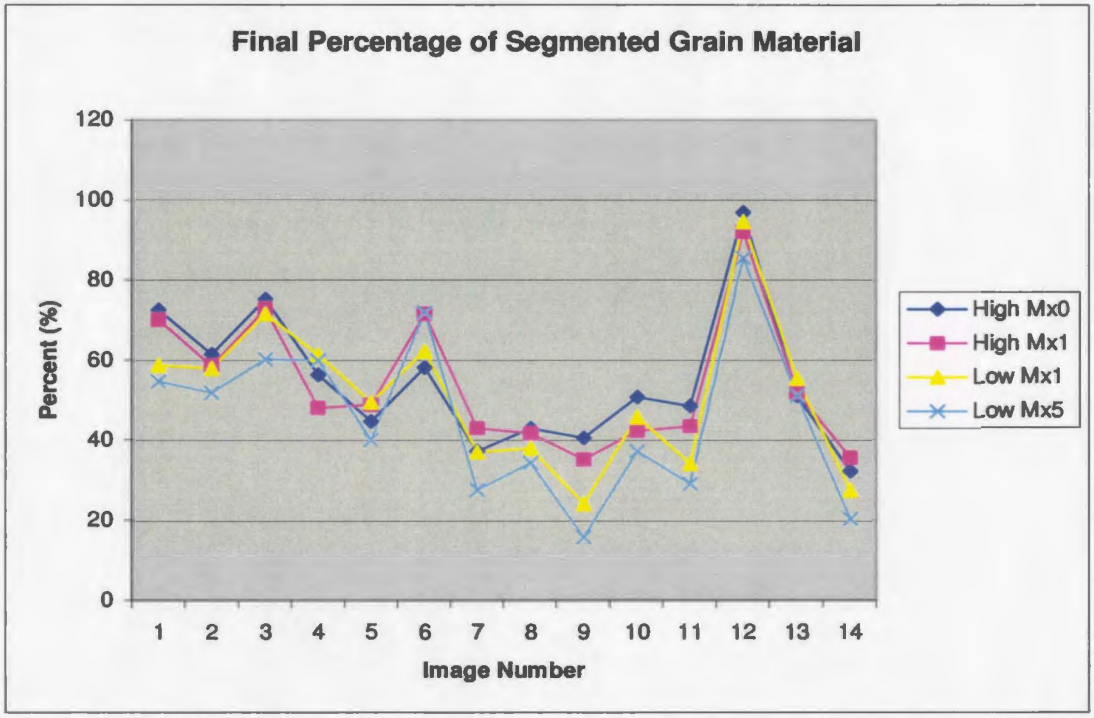

Figure 6.14: Final percentage of segmented grain material for each of the group A images.

The accuracy for each of the methods is shown in Figure 6.15. The values produced by each of the segmentation routines are similar for some images but vary significantly for others.

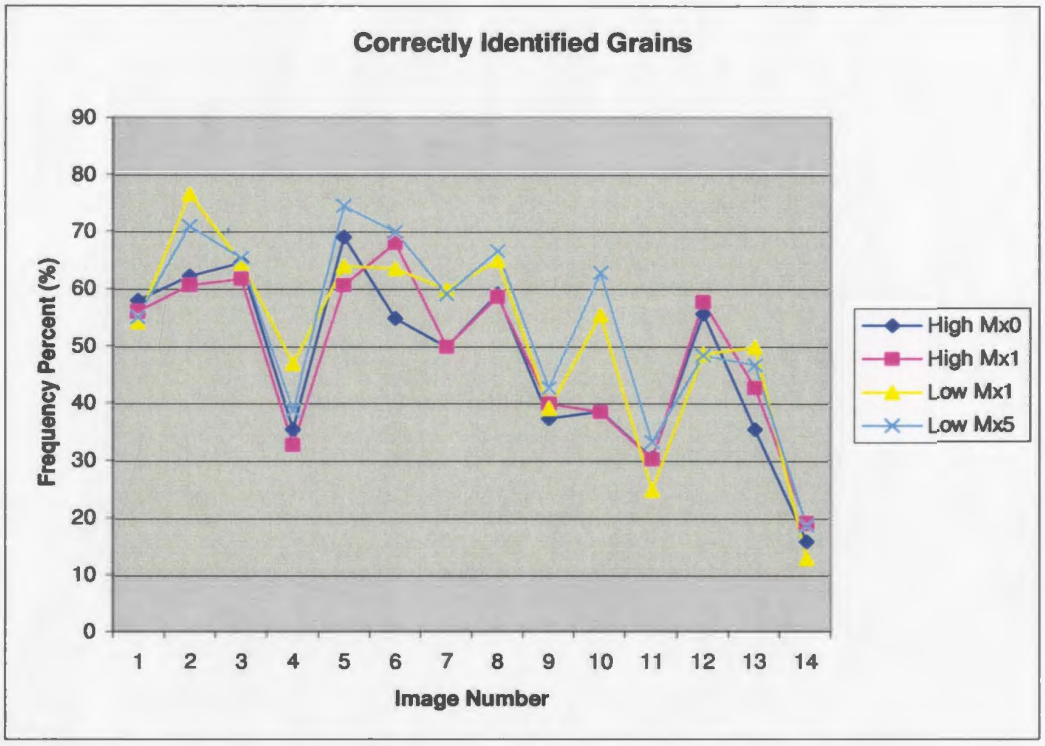

Figure 6.15: Correctly identified grains (frequency percent)

The results for the average grain area are shown in Figure 6.16. The automatically determined results for each case are compared to results from the manual segmentation. The segmentation routines produce nearly equivalent results for most images. Many of the images (1-5, 8, 10) show good agreement with the manual image results, while others (6, 7) are more spread out but are bracketed by the 'all types' value and the type 1&2 value. The remaining results are not as good with the error for images 11 and 14 being particularly high. Both of these images are composed of large grains with significant amounts of missing boundaries. The decrease in average area results from the inability to detect larger grains coupled with the fact that some grains possess texture that produces a multitude of small partial grains.

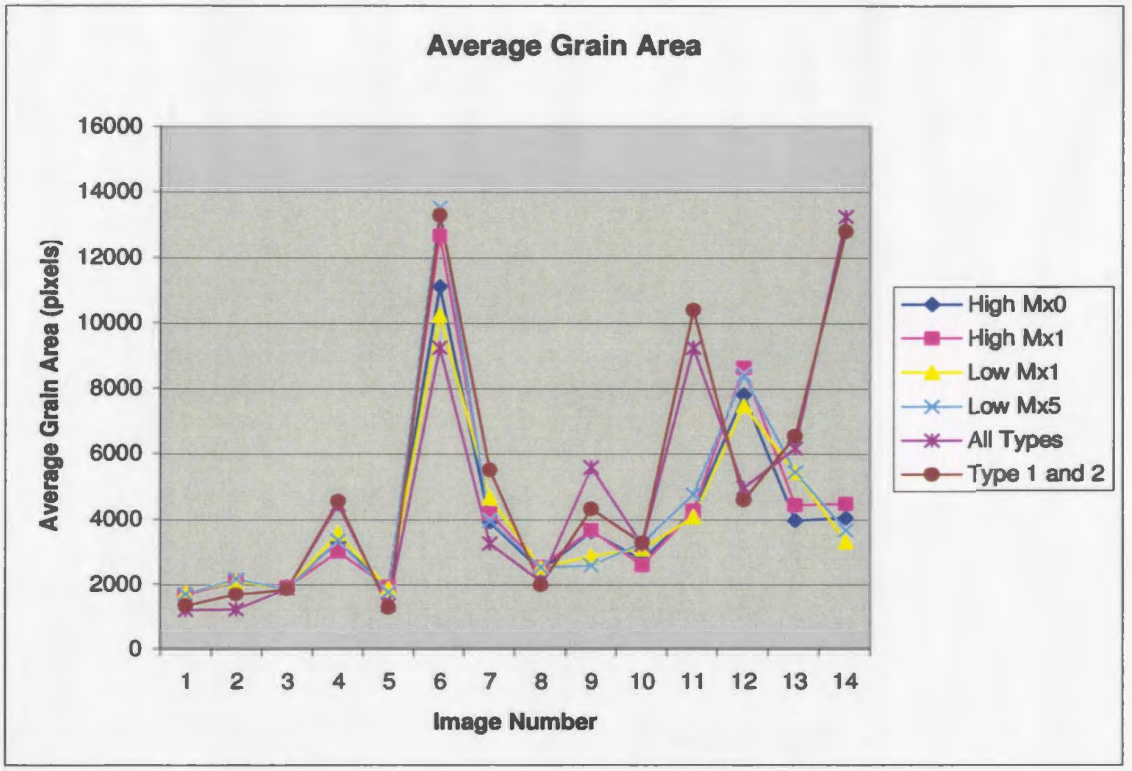

Figure 6.16: Automatic average grain area compared to manual results.

The sorting values are compared to the manual results in Figure 6.17. As with the average area values, the sorting values for each segmentation routine are very close. The sorting values show good agreement with the manual values for some images (4, 9, 12, 13), with other images (6, 7) having results that are bracketed by the manual values. The other 8 images show significant error between automatic and manual sorting values.

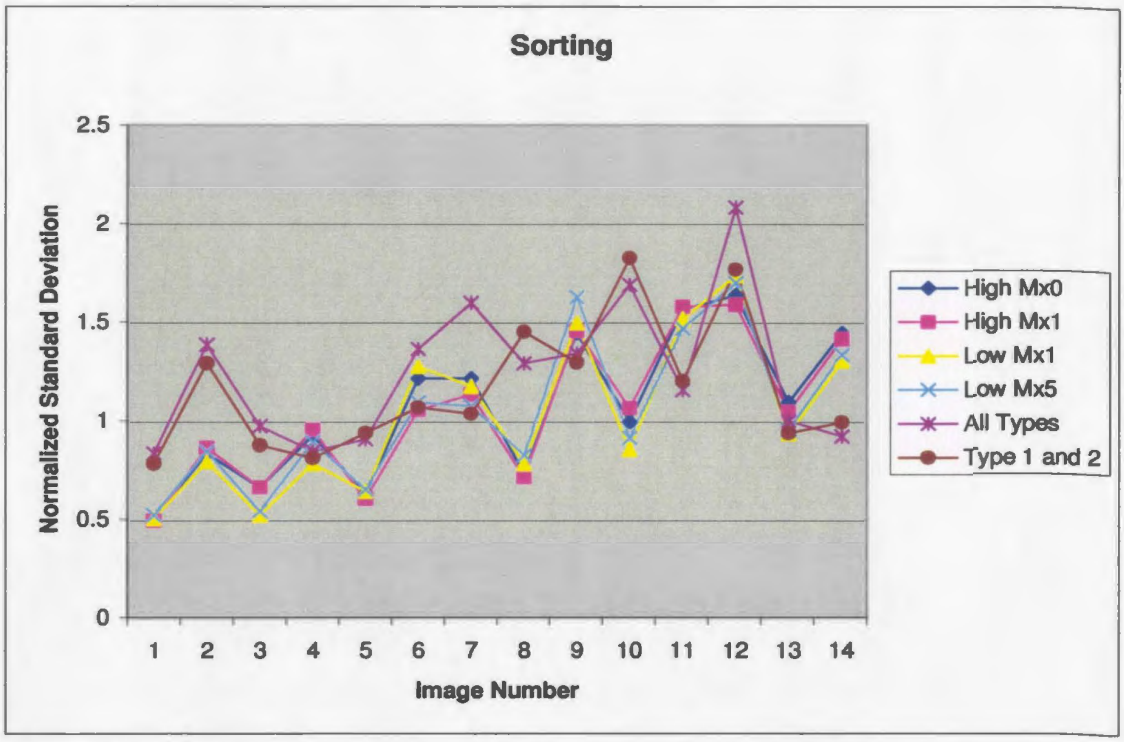

Figure 6.17: Automatic sorting measures compared to manual results.

6.4 Group B Results

Segmentation results for the B group images showed overall poor results as compared to the A group, especially for cases with very large grains. For this reason, a complete analysis of these results is not presented here, however, a few examples are given to indicate the overall results. Figure 6.18 shows an image with significant textural characteristics, mostly in the form of matrix material, and complex grain boundary interactions. The resulting segmentation is shown in Figure 6.19. The algorithm leaves much of the grain material un-segmented but it does an excellent job of classifying the high texture, indicated by pink.

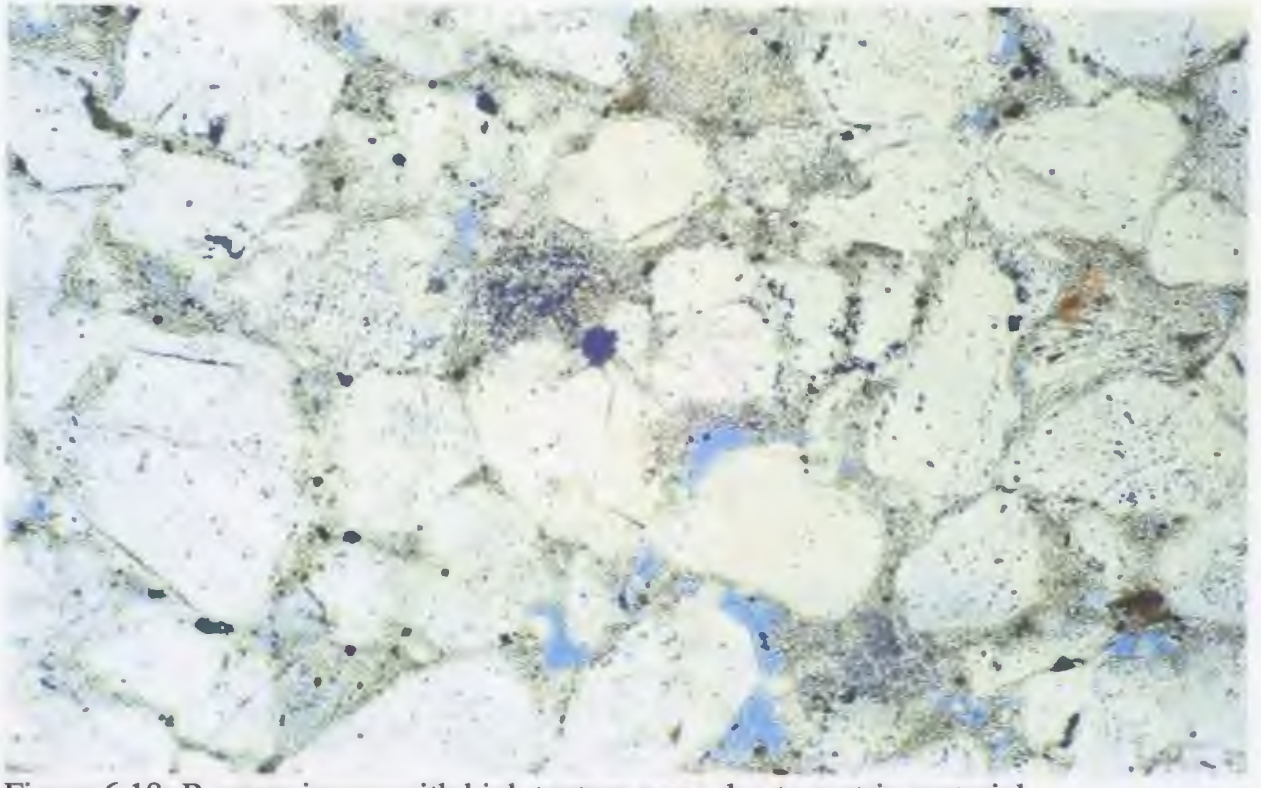

Figure 6.18: B group image with high texture areas due to matrix material.

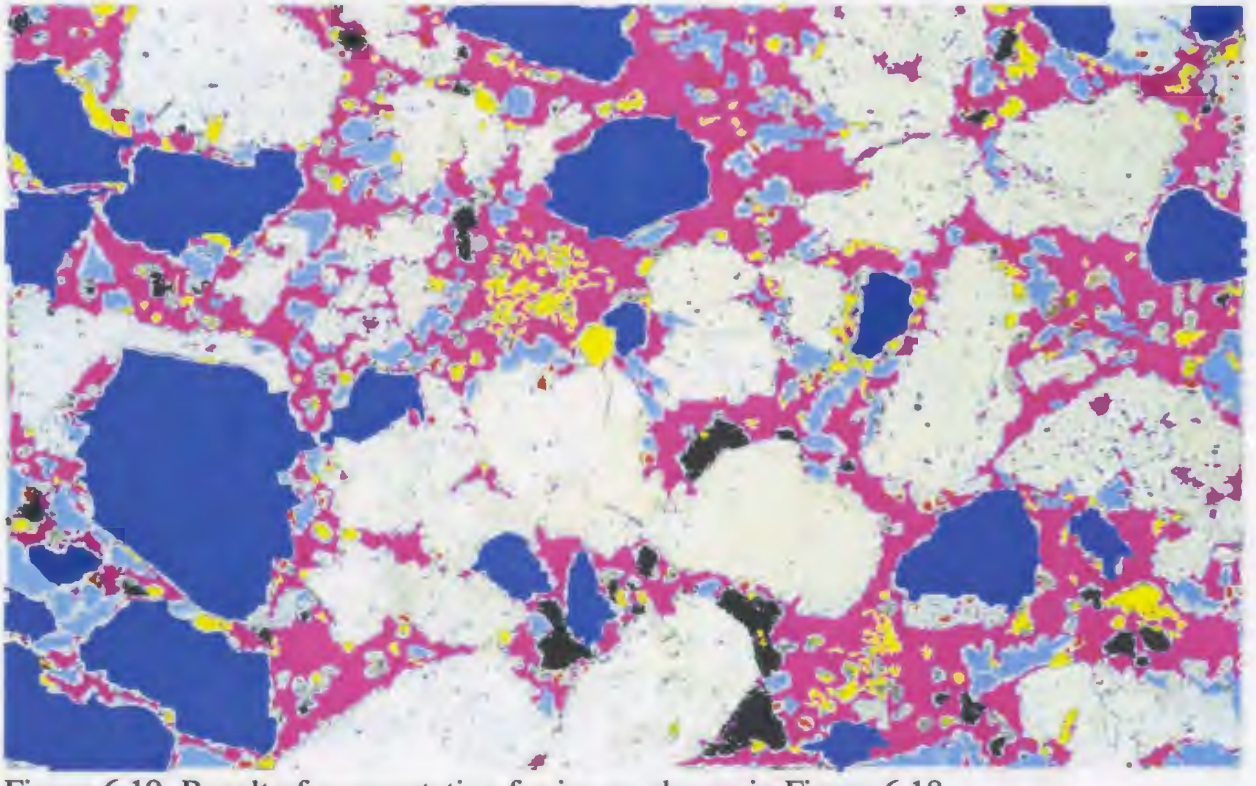

Figure 6.19: Result of segmentation for image shown in Figure 6.18.

Figure 6.20 shows another image with high texture areas but this time it is in the form of grain material. The result of the segmentation, shown in Figure 6.21, is similar to the previous example. Most grains remain un-segmented but high texture grain material is successfully identified and is shown as a combination of silt-sized grains, dirt material and high texture material. These types of results are acceptable for the B group images since they are not intended for grain size distribution. They are acquired to show fine detail such as high texture grains and matrix material and it is evident from the example that the segmentation algorithm does a good job of finding this.

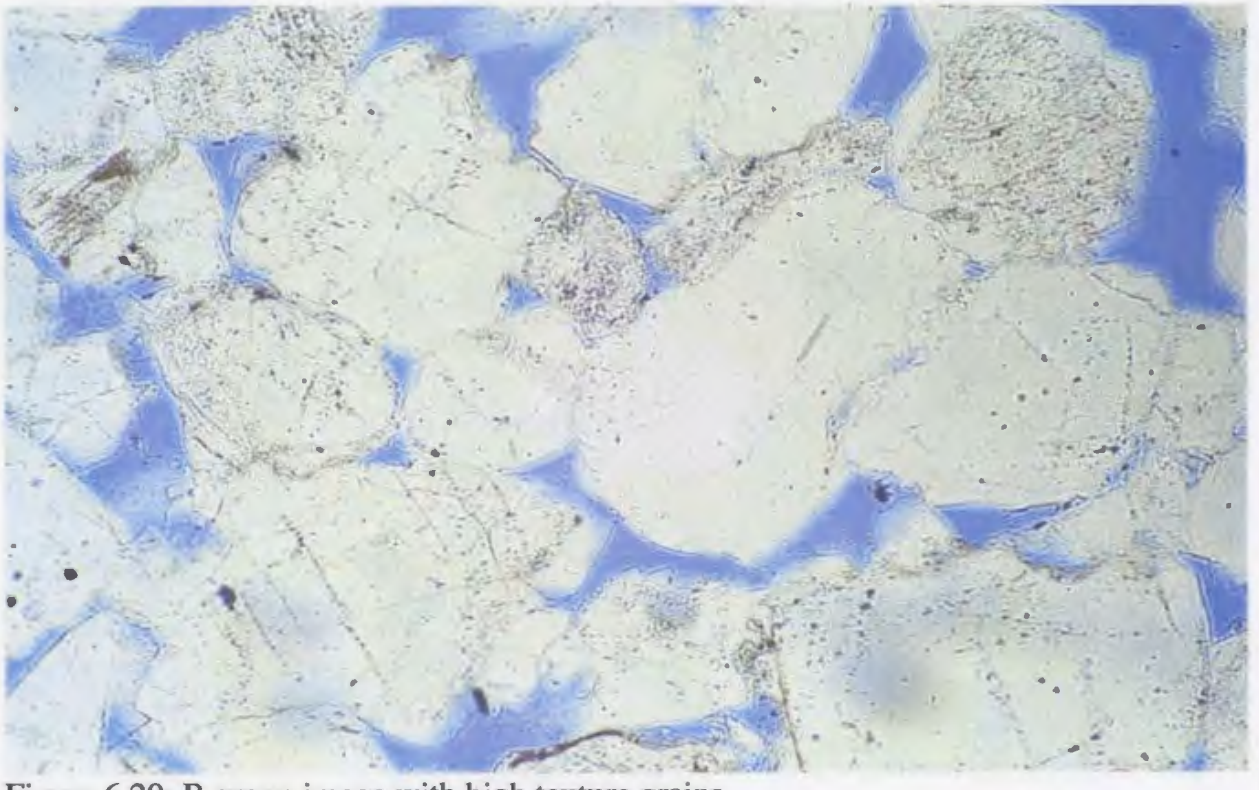

Figure 6.20: B group image with high texture grains.

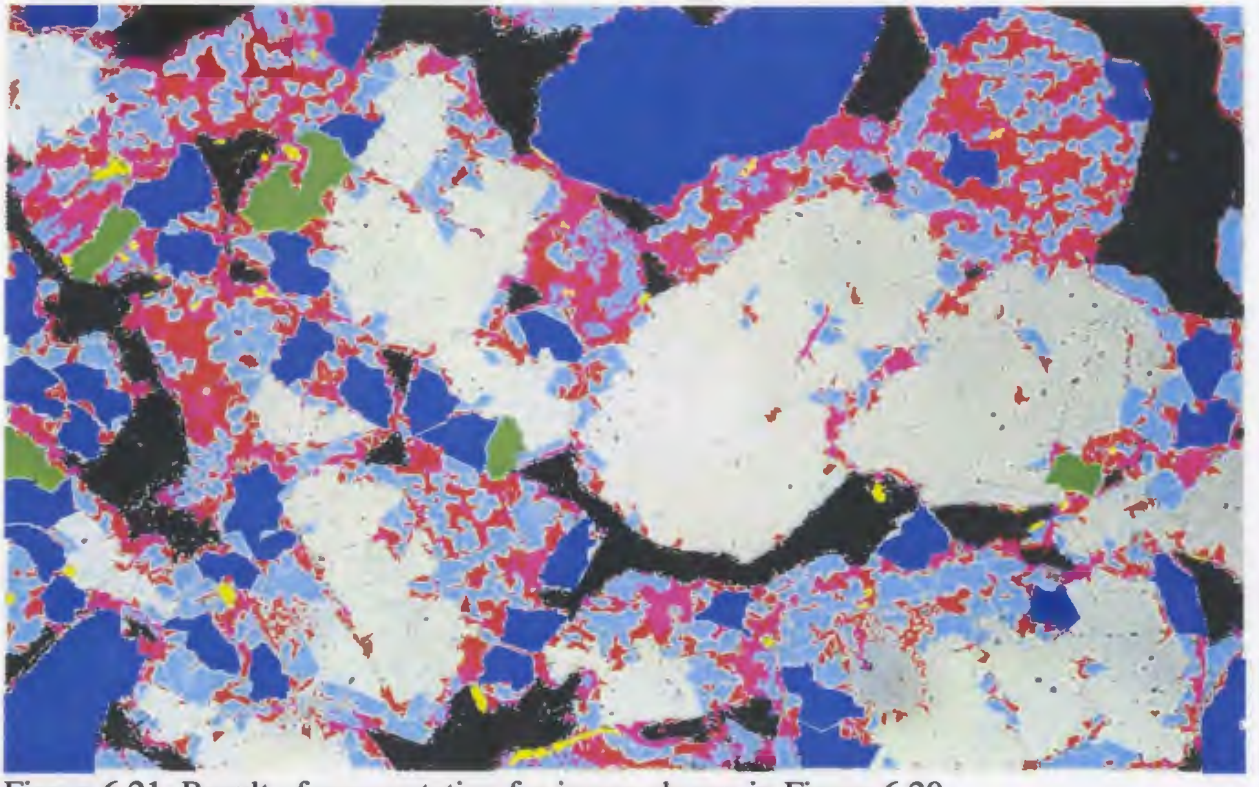

Figure 6.21: Result of segmentation for image shown in Figure 6.20.

6.5 Program Performance and Adaptability

A number of statistics are calculated as part of the output of the segmentation routine. The amount of silt-sized material is included as is the amount of "dirt" or clay-sized particles. The amount of opaque material and porosity is also measured based on simple threshholding and pixel counting operations. The percentage of image that is composed of high texture material is also included based on the Sobel texture analysis. The percentages of porosity, opaque material and high texture material are measured before the start of the segmentation process and these values do not change as the algorithm progresses. In theory, these values should be linked to other parameters such as average grain size and grain size distribution and it may be possible to select applicable

segmentation routines based on the measure. These types of relationships were explored but there were no solid links found.

Another approach is to base segmentation on the presence of one or more of the petrographic components that is produced by the segmentation process. For example, an abundance of silt-sized grains or high texture material may suggest the presence of diagenetic material of polycrystalline grains. Either, these regions can be flagged for further consideration or parameters can be adjusted to account for this and the segmentation process can be rerun in an attempt to produce fewer silt-sized grains.

An attempt was also made to automate the selection of line limits used in the line growing routines. The line limits were based on the average grain area or average pore area, depending on which is more abundant in the initial stage of segmentation. Line limits are then obtained by taking a fraction of the average radius. Low limit = ¼, mid limit = 1/3 and high the limit = ½. These limits are based on the assumption that only a small fraction of the grain boundary can be approximated by a straight line. Since the radius is an average, the low limit is applied first then the middle value followed by the high value, progressing from lower to higher as did the original routine. Results of this selection routine showed promise for some images but it was not adaptable to the entire data set.

Due to the automated accuracy routine, the program has the ability to learn from a few representative manually segmented images. That is, images that represent each of the dominant structures in a set of samples can be input and optimal algorithms can be developed for each structure. Overall, this algorithm shows great potential as a means of classifying thin-section petrographic images.

Chapter 7

Conclusions and Recommendations

7.1 Conclusions

The image segmentation and measurement algorithms presented here were applied to a data set containing 28 images. The goal was to partially automate the analysis performed by an experienced operator, thereby allowing significantly more images to be processed. For the group A images the algorithm does a good job of segmenting and identifying type 1 and type 2 quartz grains. For the group B images, grains are often too large or the magnification shows too much detail and, therefore, the segmentation routine is focused on measuring image constituents other than quartz. Major conclusions are as follows:

- It is clear that an integrated approach is necessary to accurately segment thin-section petrographic images. This is due to the fact that the grain types, and hence the grain contacts, vary significantly throughout the images.
- The effect of operator error is noted in other research and despite efforts to minimize this, it is still a factor here. In several cases, the segmentation routine correctly identifies grains that are incorrectly identified by manual analysis.
- Accurate segmentation is highly dependent on grain size with larger grains, usually found in group B, being more difficult to segment due the fact that longer boundaries need to be approximated. However, due to

their size, large grains are often cropped and therefore their segmentation is meaningless since the entire grain is not present to be measured.

- Diagenetic effects, in the form of quartz overgrowths, compaction, including shear bands and fracturing, dissolved or partially dissolved grains and authigenic clay material, play an important part in the segmentation process.
- Polycrystalline quartz grains, quartz with abundant inclusions, and feldspars are noted throughout the samples and these types of grains (type 3) complicate the segmentation by creating partial grains.
- Grain fabric, left un-segmented at the conclusion of the algorithm, is typically composed of very large grains, areas of significant diagenetic alteration, or a combination of both. Attention can be drawn to images with large amounts of un-segmented fabric and they can either be quickly segmented by the operator or they can be studied for their diagenetic properties; either way, it is important to note these features.

While conducting this research, it became apparent that many of the grains are altered through the process of diagenesis and some means to quantify this effect would prove useful. In fact, the petrographic report completed for these samples [El-Dein et al., 1984] shows that only a small part of the analysis conducted by a petrographic laboratory is focused on grain area and sorting. Significantly more consideration is given to other components, that better characterize the sediment. This includes the presence of any diagenetic material including authigenic clays and altered quartz or feldspar grains.

An attempt is made to lay a solid foundation upon which additional methods and analysis can be based. This included a significant amount of time and effort spent to understand the process of characterizing sediments using thin-section imaging. This effort has paid off and produced a well-developed analysis of the problems associated with using thin-section analysis to produce data for characterization of geological structures. It was important to consider the use of these data throughout the development of the algorithms.

7.2 Recommendations

The groundwork developed in this thesis gives valuable insight into the automated segmentation of thin-section petrographic images. Many recommendations can be made based on this work; the most significant of these are as follows:

- Additional blob reconstruction strategies can be developed to reduce the effect of over-segmented grains. These strategies would likely include texture analysis.
- Under-segmented grains could benefit from additional analysis at the grain recognition stage. For example, many of the under-segmented grains meet the splitting criteria as defined by the modified watershed transform developed in previous work [Zhao, 2000].
- Further investigation into porosity and rock type relationships could help in the automatic characterization process. If the program could identify distinct rock types based on porosity analysis, the operator could segment characteristic images, input them into the program, and the optimal segmentation routine for that particular rock type can be found.

- Work with a petrographer to produce optimal images for automated analysis.

References

Bowers, Mark C., Ehrlich, Robert, Clark, Ross A. (1994). Determination of Petrographic Factors Controlling Permeability Using Petrographic Image Analysis and Core Data, Satun Field, Pattani Basin, Gulf of Thailand, Marine and Petroleum Geology Volume 11, Number 2.

Casadei, S., Mitter, S. (1998). Hierarchical Image Segmentation – Part I: Detection of Regular Curves in a Vector Graph. International Journal of Computer Vision 27(1), 71-100.

Clelland, W.D., Fens, T.W. (1991). Automated Rock Characterization with SEM/Image-Analysis Techniques. SPE Formation Evaluation, December, 437-443.

Crabtree, S.J., Ehrlich, R., Prince, C. (1984). Evaluation of strategies for segmentation of blue-dyed pores in thin sections of reservoir rocks. Computer Vision, Graphics and Image Processing 28, 1-18.

Davies, D.K. (1990). Image Analysis of Reservoir Pore Systems: State of the Art in Solving Problems Related to Reservoir Quality. SPE Formation Damage Control Symposium, February 1990.

DeLesse, M.A. (1847). Procede mechanique pour determiner la composition des roches. C.R. Acad. Sci., Paris 25, 544.

Diogenes, A.N., Haendchen, C., Hoff, E., Maliska Jr., C.R., Fernandes, C.P., Maliska, A.M. (2003). Low Carbon Steel Grains Image Analysis Using Microstructural Characterization Software Imago. Acta Microscopica, 12, Supplement C. 267-268.

Ehrlich, R., Crabtree, S. J., Horkowitz, K. O., and Horkowitz, J. P. (1991a). Petrography and Reservoir Physics I: Objective Classification of Reservoir Porosity, The American Association of Petroleum Geologists Bulletin, 75, 1579-1592.

Ehrlich, R., Etris, E. L., Brumfield, D., Yuan, L. P. and Crabtree, S.J. (1991b). Petrography and Reservoir Physics III: Physical Models of Permeability and Formation Factor, The American Association of Petroleum Geologists Bulletin, 75, 1547-1562.

El-Dein, S.S., Sitek, J., Ibrahim, H., Davies, G.R. (1984). Petrographic Analysis of Mobil et al Hibernia B-27 Offshore Eastern Canada, Volume I. AGAT Technologies Inc. 3650 – 21st. Street N.E., Calgary, Alberta, T2E 6V6.

Francus, P. (1998). An image-analysis technique to measure grain-size variation in thin sections of soft clastic sediments. Sedimentary Geology 121, 289-298.

Gies, R.M., McGovern J. (1993). Petographic Image Analysis: An Effective Technology for Delineating Reservoir Quality. SPE Gas Technology Symposium, June 1993.

Gosine, Ray and Burden, Elliott (1999), PetrograFX: A petrographic Image Analysis and Database Management Toolkit, A Proposal to Seed Collaborative Reserch in Information Technology and Informatics, C-CORE & Earth Sciences, Memorial University of Newfoundland.

Gonzalez, R. C., Woods, R.E. (2002) Digital Image Processing, Addison-Wesley Publishing Company.

Gribble, Colin D., Hall, Allan J. (1992) Optical Mineralogy: Principles and Practice, UCL Press Limited, London.

Hibernia. HMDC_web_page\Core_data\Petrographic_data\Petrographic_web_pages\ Hibernia_B-16_2\Hib_B16-2_pet_inventory.html.

James, Robert A. (1995). Application of Petrographic Image Analysis to the Characterization of Fluid-Flow Pathways in a Highly-Cemented Reservoir: Kane Field, Pennsylvania, U.S.A., Journal of Petroleum Science & Engineering 13, 141-154.

Kennedy, Stephen K. and Mazzullo. Jim (1991), Image Analysis Method of Grain Size Measurement, In Principles, Methods and Application of Particle Size Analysis, Syvitski, J.P.M (ed). pp 76-87.

Koederitz, L. F., Harvey, A. H., Honarpour, M. (1989) Introduction to Petroleum Reservoir Analysis, Gulf Publishing Company, Houston.

Kraatz, M., Heliot, C., Blum, W., Zschech, E. (2003). Automated Grain Analysis in the Sub-Micron Range with AFM Image Processing. Prakt. Metallogr. 40, 69-77.

Lester, E., Allen, M., Cloke, M., Miles, N. J. (1993). Image Analysis Techniques for Petrographic Analysis, Fuel Processing Technology, 36, 17-24.

Lewis, Douglas W., McConchie, David. (1994a) Analytical Sedimentology, Chapman & Hall, New York.

Lewis, Douglas W., McConchie, David. (1994b) Practical Sedimentology, Chapman & Hall, New York.

Mahadevan, S., Casasent, D. (2003). Automated image processing for grain boundary analysis. Ultramicroscopy 96, 153-162.

Malik, J., Belongie, S., Leung, T., Shi, J. (2001). Contour and texture analysis for image segmentation. International Journal of Computer Vision 43(1), 7-27.

Matrox Imaging Library User's Guide (1999). Version 6.0. Matrox Electronic Systems Ltd..

McCreesh, C. A., Ehrlich, R. and Crabtree, S.J. (1991). Petrography and Reservoir Physics II: Relating Thin Section Porosity to Capillary Pressure, the Association Between Pore Types and Throat Size, The American Association of Petroleum Geologists Bulletin, 75, 1563-1578.

Mowers, Theodore T., Budd, David A. (1996). Quantification of Porosity and Permeability Reduction Due to Calcite Cementation Using Computer-Assisted Petrographic Image Analysis Techniques, The American Association of Petroleum Geologists Bulletin, V. 80, No. 3, 309-322.

Radaelli, F., Balzarini M., Nicula S., Ortenzi, A. (1998). Rock Structure Characterization Through Imaging Techniques Integration, SPE European Petroleum Conference, October, 1998.

Raymond, Loren A. (2002) The Study of Igneous, Sedimentary & Metamorphic Rocks, Second Edition, McGraw-Hill, New York.

Ross, B.J., Fueten, F., Yashkir, D.Y. (2001). Automatic mineral identification using genetic programming. Machine Vision and Applications 13, 61-69.

Russ, John, C. (1990) Computer-Assisted Microscopy: The Measurement and Analysis of Images, Plenum Press, New York.

Russ, John, C. (1995) The Image Processing Handbook, CRC Press, Boca Raton.

Ruzyla, K. (1986). Characterization of Pore Space by Quantitative Image Analysis. SPE Formation Evaluation, August 1986.

Scholle, Peter A. (1979). A Color Illustrated Guide To Constituents, Textures, Cements, and Porosities of Sandstones and Associated Rocks, Memoir 28, The American Association of Petroleum Geologists.

Sippel, R.F. (1968). Sandstone Petrology, Evidence from Luminescence Petrography. Journal of Sedimentary Petrology, 38, 530-554.

Smith, S.M., Brady, J.M. (1997). SUSAN – A New Approach to Low Level Image Processing. International Journal of Computer Vision, 23(1), 45-78.

Starkey, J., Samantaray, A.K. (1994). A Microcomputer-based System for quantitative petrographic analysis. Computers and Geosciences 20, 1285-1296.

Thomson, R.C., (1996). Micrograph Analysis and image understanding. Proceedings from the IEEE southwest symposium on image Analysis and Interpretation.

Tomutsa, L., Mahmood, S.M., Brinkmeyer A., Honarpour M. (1990). Application of Integrated Pore-to-Core Image Analysis to Study Fluid Distribution in Reservoir Rocks. 65th Technical Conference and Exhibition of the Society of Petroleum Engineers, September 1990.

Tucker, Maurice E. (2001). Sedimentary Petrology, Third Edition, Blackwell Science Ltd.

Van den Berg, E.H., Meesters, A.G.C.A., Kenter, J.A.M., Schlager, W. (2002). Automated separation of touching grains in digital images of thin sections. Computers and Geosciences 28, 179-190.

Zhao, Xianghong (2000). Automated Image Analysis for Petrographic Images Assessments. M.Eng., Memorial University of Newfoundland.