

**MEASURING BORDER IRREGULARITY AND SHAPE
OF CUTANEOUS MELANOCYTIC LESIONS**

by

Tim Kam Lee

B.Sc., The University of British Columbia, 1980

M.Sc., The University of British Columbia, 1983

THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

in the School
of
Computing Science

© Tim Kam Lee 2001

SIMON FRASER UNIVERSITY

January 2001

All rights reserved. This work may not be
reproduced in whole or in part, by photocopy
or other means, without permission of the author.

APPROVAL

Name: Tim Kam Lee
Degree: Doctor of Philosophy
Title of thesis: Measuring Border Irregularity and Shape of Cutaneous Melanocytic Lesions

Examining Committee:

Chair: Dr. Arvind Gupta

Dr. M. Stella Atkins, Senior Supervisor

Dr. Ze-Nian Li, Supervisor

Dr. Calum MacAulay, Supervisor
Dept. of Pathology and Laboratory Medicine
University of British Columbia

Dr. David McLean, Supervisor
Division of Dermatology
University of British Columbia

Dr. Mark Drew, SFU Examiner

Dr. Ela Claridge, External Examiner
School of Computer Science
University of Birmingham

Date Approved: _____

Abstract

Cutaneous melanocytic lesions, commonly known as moles, are mostly benign; however, some of them are malignant melanomas, the most fatal form of skin cancer. Because the survival rate of melanoma is inversely proportional to the thickness of the tumor, early detection is vital to the treatment process. Many dermatologists have advocated the development of computer-aided diagnosis systems for early detection of melanoma.

One of the important clinical features differentiating benign nevi from malignant melanomas is the lesion border irregularity. There are two types of border irregularity: texture and structure irregularities. Texture irregularities are the small variations along the border, while structure irregularities are the global indentations and protrusions that may suggest either the unstable growth in a lesion or regression of a melanoma. An accurate measurement of structure irregularities is essential to detect the malignancy of melanoma.

This thesis extends the classic curvature scale-space filtering technique to locate all structure irregular segments along a melanocytic lesion border. An area-based index, called irregularity index, is then computed for each segment. From the individual irregularity index, two important new measures, the most significant irregularity index and the overall irregularity index, are derived. These two indices describe the degree of irregularity along the lesion border .

A double-blind user study is performed to compare the new measures with twenty experienced dermatologists' evaluations. Forty melanocytic lesion images were selected and their borders were extracted automatically after dark thick hairs were removed by a preprocessor called DullRazor. The overall irregularity index and the most significant irregularity index were calculated together with three other common shape descriptors. All computed measures and the dermatologists' evaluations were analysed statistically. The results showed that the overall irregularity index was the best predictor for the clinical evaluation, and both the overall irregularity index and the most significant irregularity index outperformed the other shape descriptors. The new method has great potential for computer-aided diagnosis systems.

Dedication

*To my wife Chao-Yong and
my children Jenine and Jensen*

Acknowledgments

I wish to express my deepest gratitude to my senior supervisor Stella Atkins. She patiently guided me through the course of my study program. Without her encouragement and enthusiasm, the completion of this dissertation would not been possible.

I am deeply indebted to my supervisor David McLean, who have been continuously supporting my research in skin cancer by sharing his medical knowledge and scientific research methodologies. He was instrumental in completing the research project.

I am grateful to my supervisors Ze-Nian Li and Calum MacAulay, my external examiner Ela Claridge and SFU examiner Mark Drew for their valuable comments and advice, which greatly improved the dissertation and broadened my research interest.

Special thanks are reserved for Vincent Ng who introduced me to the area of medical image analysis. Our early collaboration built a strong foundation for my research program.

My appreciation goes to Chris Bajdik, Andy Coldman, Nhu Le and Richard Lee for their constructive discussions and suggestions to my research project, which greatly influenced the final shape of the dissertation.

After working for 12 years, I was most reluctant to return to school. Without the encouragement and support from Pierre Band, Michele Deschamps, Brian Funt, and Rick Gallagher, I would not have taken the initial step of enrolling into the program and have completed my study while holding a full-time job. I would like to express my sincere thanks to all of them.

I would like to thank all the colleagues and friends in Cancer Control Research Program, BC Cancer Agency, Centre for Systems Science and School of Computing Science. Their friendship made my work and study environment enjoyable even when I attempted to cope between work and study.

I would like to acknowledge BC Health Research Foundation for providing funding (grants #182(92-2) and #142(97-2)) to support my research program.

Finally, I am most thankful to my wife Chao-Yong for her unlimited love, care and encouragement. I always amazed how she could find energy and time to support and contribute in every aspect of my research work, after a full day at work and a busy evening with the children. Being a graduate from Computer Science, she fully appreciated my research problems and she was always the first person to hear my ideas. My children Jenine and Jensen also deserved my appreciation. Their music preserved my sanity and I learned the positive meaning of life from their youthful and bright eyes.

Table of Contents

Approval	ii
Abstract	iii
Dedication.....	iv
Acknowledgments	v
Table of Contents	vii
List of Tables.....	xi
List of Figures	xii
1. Introduction	1
1.1. Motivation	1
1.2. Melanocytic lesions	1
1.2.1. Benign nevi	3
1.2.2. Malignant melanomas	4
1.3. Public health programs	6
1.4. Non-invasive computer-aided diagnostic systems	7
1.4.1. Image acquisition.....	8
1.4.2. Image processing	11
1.4.3. Lesion classification.....	11
1.5. Border shape as a defining feature	12
1.6. Outline of the thesis	14
2. Related Work	15
2.1. Published border shape descriptors for melanocytic lesions	15
2.1.1. Global measures based on a single-scale	15
2.1.2. Global measures based on multi-scale	16
2.2. Other shape algorithms in computer vision.....	19
2.2.1. Parameterization of object borders	20

2.2.2. Shape descriptors	20
2.2.2.1. Fourier Descriptors	21
2.2.2.2. Moments analysis	22
2.2.3. Object partitioning.....	24
2.2.3.1. Skeleton based techniques.....	24
2.2.3.2. Boundary based techniques	26
2.2.4. Multi-scale methods	28
2.2.4.1. Wavelet transform.....	29
2.2.4.2. Scale-space filtering	30
2.3. Multimedia shape coding.....	33
2.3.1. Vertex-based shape coding	33
2.3.2. Baseline based shape coding	34
2.3.3. Modified Modified Reed shape coding	34
2.3.4. Context-based arithmetic encoder	34
3. Data Collection and Preprocessing	36
3.1. Data collection.....	36
3.1.1. Images from a video microscopy device.....	36
3.1.2. Images from 35-mm colour slides	38
3.2. Preprocessing	38
3.2.1. DullRazor.....	38
3.2.2. Segmentation	44
3.2.3. Issues arising in the data preprocessing phase.....	50
4. Irregularity Indices	54
4.1. Abstracting the lesion border	54
4.2. Defining indentation and protrusion segments.....	55
4.3. Extracting structure indentation and protrusion segments.....	57
4.3.1. Scale dependence of irregularity segments.....	57

4.3.2. Classic curvature scale-space	58
4.3.3. The sigma-ratio shape measure.....	61
4.3.4. Extended curvature scale-space.....	63
4.3.5. Comparison of classic and extended curvature scale-space images	67
4.3.6. Identifying the span of indentation/protrusion segments	73
4.3.7. Hierarchical structures for indentation/protrusion segments.....	74
4.3.8. Flat irregularity segments	76
4.4. Calculating irregularity indices	77
4.5. Sensitivity analysis of irregularity indices	80
4.6. The advantages of irregularity indices	81
5. A User Study	83
5.1. Method.....	83
5.1.1 Gold standard	83
5.1.2. Assembling the data	83
5.1.3. The compactness index implementation.....	86
5.1.4. The fractal dimension implementation.....	86
5.1.5. The irregularity index implementation.....	90
5.1.6. Clinical evaluation.....	90
5.1.7. Statistical methods.....	91
5.2. Results.....	93
5.3. Discussion.....	93
5.3.1. Clinical evaluation and gold standard.....	93
5.3.2. Relationship between automatic methods and the average clinical evaluations	98
5.3.3. A single point estimation vs. a measurement set	99
5.3.4. Linear regression model.....	100

5.3.4.1. Single group analysis	100
5.3.4.2. Subgroup analysis	101
5.3.5. Overall irregularity index and histology	103
5.3.6. Summary	105
6. Conclusion and Future Work	107
6.1. Contribution and originality	107
6.2. Future work	110
Appendix A	112
References	118

List of Tables

1.1. Melanoma thickness and 10 year survival rates.....	6
4.1. Comparison of the SR, OII and MSII for two phantoms	81
5.1. Clinical evaluation results of the lesion borders shown in Fig. 5.2.....	94
5.2. Average clinical evaluation and tested measurements	95
5.3. Spearman coefficients for all test data.....	97
5.4. Multiple linear regression analysis results.....	98

List of Figures

1.1. Canadian age-standardized incidence rates for melanomas	2
1.2. Skin structure	3
1.3. Melanocytic lesions	4
1.4. Schematic diagram for a non-invasive computer-aided diagnostic system.....	8
1.5. Obtaining subsurface view using two polarized filters	10
1.6. Subsurface view of a melanocytic lesion.....	11
1.7. Three lesions and their borders	13
2.1. Fractal dimension and structure fractal dimension.....	18
2.2. Computing the SAT by mathematical morphology method.....	25
2.3. Limb-based and neck-based part-lines	28
2.4. Scale-space filtering for a 1D function.....	31
2.5. Scale-space filtering for a closed contour	32
3.1. The video microscopy imaging device.....	37
3.2. A melanocytic lesion image covered by dark thick hairs	39
3.3. Structure element for the generalized closing operation.....	40
3.4. Hair mask M after applying generalized closing operation to Fig. 3.2.....	41
3.5. The cleanup hair mask of Fig. 3.4	42
3.6. The skin image of Fig. 3.2 after the replacement step.....	43
3.7. The enlarged hair mask for Fig. 3.5	43
3.8. The final output of DullRazor for Fig. 3.2	44
3.9. Filter windows for median filter.....	45
3.10. Scheme for the three-stage CWMF.....	46
3.11. Effects of the three-stage CWMF on Fig. 3.8	47
3.12. The red, green and blue band of Fig. 3.8.....	49
3.13. Segmentation result	50

3.14. A skin lesion heavily covered by dark thick hairs	51
3.15. Example of good segmentation results	52
3.16. Example of fair segmentation results.....	53
3.17. Examples of poor segmentation results	53
4.1. Example of an extracted lesion border.....	55
4.2. Definition for indentations and protrusions	56
4.3. Gaussian smoothing process for lesion border shown in Fig. 4.1	60
4.4. Classic curvature scale-space image	61
4.5. Curvature function of the smoothed border at $\sigma = 40$ of Fig. 4.3.....	62
4.6. Three lesion borders of Fig. 1.7.....	62
4.7. Two phantoms for sigma-ratio	63
4.8. The extended and curvature scale-space images.....	65
4.9. Curvature function and its first derivative for $\sigma = 48$ of Fig. 4.3.....	66
4.10. Smoothed borders of Fig. 4.3 at $\sigma = 40$ and $\sigma = 63$	68
4.11. Smoothed borders of Fig. 4.3 at $\sigma = 48$ and $\sigma = 101$	70
4.12. Hierarchical relationship between segments	75
4.13. Smoothing effect on indentation and protrusion.....	78
4.14. Indentation/protrusion segments of Fig. 4.1	80
4.15. The most significant indentations for two phantoms shown in Fig. 4.7.....	81
5.1. The overall irregularity index (OII) of 108 lesions	84
5.2. Forty lesion borders selected for the user study	85
5.3. An example of the log-log plot for fractal dimensions	87
5.4. Koch snowflake and its construction procedure	88
5.5. Koch square flake and its construction procedure	89
5.6. The most significant indentation/protrusion segments	96
5.7. Plotting the tested measurements vs. clinical evaluation.....	97
5.8. The re-plot of Fig. 5.7.....	102

5.9. Predicting average clinical evaluation by OIL.....	103
6.1. A skin image covered by light coloured hairs	110

Chapter 1

Introduction

1.1. Motivation

The most fatal form of skin cancer, cutaneous malignant melanoma, has been increasing rapidly in the western world for the last few decades [7, 45, 64, 98]. Canada has an incidence rate that reflects the world trend [107] (see Fig. 1.1) and melanoma has become the third most common cancer for people between 30 and 40 years of age [112]. As the diagnosis frequently occurs at young adulthood when the family and financial commitment is the greatest, the potential person-years of life loss has a tremendous impact on the individual and his/her family. In order to reduce the mortality rate, which is proportional to the thickness of the tumor, many dermatologists have advocated the development of computer-aided diagnosis systems for early detection of melanoma. The development of such diagnostic devices begins with the understanding of the benign form (nevus) and the malignant form (melanoma) of a melanocytic lesion.

1.2. Melanocytic lesions

Melanocytic lesions are the brown pigmented skin lesions commonly known as moles. They are formed by nests of specialized cells called melanocytes. Normal melanocytes live in isolation as solitary units in the bottom layer, stratum basale, of the outmost part of the skin, the epidermis¹ [49, 144, 155]. (See Fig. 1.2) Intermixing with basal cells (keratinocytes), melanocytes make

¹ Human skin consists of three principal parts: epidermis, dermis and subcutis. The epidermis has four to five layers of cells. Acting as the protective barrier of the body, the basal cells, keratinocytes, continuously migrate upward from stratum basale, the bottom layer of the epidermis. During the migration, these cells are flattened and mutate at the upper layers. Finally, the dehydrated dead cells are shed off the body. The entire migration and maturation processes take about a month. The dermis is the main part of the skin structure, which is separated from the epidermis by a basement membrane, the dermoepidermal junction. There are dense connective tissues, blood vessels, sweat glands, hair follicles and nerve endings in the dermis. Inferior to the dermis is subcutis, which consists of loose connective tissues and fat.




Figure 1.1. Canadian age-standardized incidence rates (ASIR) for melanomas. The rates for 1996 to 2000 are estimates. Source: National Cancer Institute of Canada: Canadian Cancer Statistics 2000, Toronto, Canada, 2000.

up about 5% to 10% of the cell population in the basal layer and they synthesize a brown pigment called melanin, which is transferred to nearby keratinocytes via dendritic processes. As keratinocytes move upwards during their migration journey, melanin is transported to the upper layers of the epidermis [55, 128]. Melanin absorbs a broad range of the optical spectrum from 250 nm to 1200 nm [1, 3, 70]. The absorption is not uniform for the entire spectrum. The strongest attenuation occurs at the shortest wavelengths, especially at the ultraviolet (UV) wavelength, and there is minimum absorption below 1200 nm. When melanocytes are irradiated by sunlight, melanin absorbs the radiation energy and becomes darker in colour [49, 144]. Prolonged radiation increases the number of melanocytes and the production of melanin. Hence, melanocytes shield the body from harmful UV radiation. When a group of melanocytes congregates together and forms a lesion, a brown pigmented patch appears on the skin due to the high concentration of melanin. These melanocytic lesions may be composed of benign or malignant cells [99].




Figure 1.2. Skin structure

1.2.1. Benign nevi

Benign melanocytic lesions (also called pigmented skin lesions or pigmented nevi) usually have a round or oval shape with a regular contour and uniform colour. Sometimes the lesion is raised [99]. Fig. 1.3a shows a typical benign nevus. The course of evolution for these lesions, from junctional nevi to compound nevi to intradermal nevi, may be revealed by counting the number of nevi in various age groups and studying lesion specimens under a microscope. About 1% of infants are born with congenital nevi [14, 57, 99]. Acquired nevi soon appear in the early stages of childhood. Most of these nevi are junctional nevi, which are small nests of melanocytes residing in the dermoepidermal junction. The number of nevi then increases rapidly to an average count of 15 to 40 in adolescence and early adulthood. The increase is positively

associated with the amount of solar radiation the individual is exposed to in an unprotected state [47, 48]. The common nevus type in the young adult age group is the compound nevus, which has part of the lesion in the dermoepidermal junction and part of the lesion migrating down to the enlarged papillary dermis. After adulthood, the number of nevi declines and pigmented nevi are rare after age 80. This phenomenon can be explained by the evolution of compound nevi to intradermal nevi: the compound nevi leave the dermoepidermal junction and descend entirely to the dermis. Intradermal nevi are usually only slightly pigmented because melanocytes seldom undergo cell division in the dermis. The number of melanocytes declines due to attrition [99].




Figure 1.3. Melanocytic lesions (a) Benign nevus (b) Malignant melanoma

1.2.2. Malignant melanomas

Malignant melanomas are usually described as enlarged lesions with multiple colour shades. Furthermore, their borders tend to be irregular and asymmetric with protrusions and indentations [99, 124]. (Fig. 1.3b shows a malignant melanoma.) The formation of malignant lesions usually begins with melanocytes damaged by UV radiation² [7, 12, 46, 124]. The damaged cells may be solitary melanocytes in the normally appearing skin (de novo) or a sub-

² Similar to benign nevi, the major risk factors for malignant melanomas are solar UV radiation and individual's genetic susceptibility.

population of the melanocytes in a pigmented lesion. Fortunately, further development of damaged melanocytes is usually stopped by the body's safeguards. Only when all the body's security checks are bypassed, can the run-away malignant cells form a cancer by becoming immortal and proliferating indefinitely [37, 147, 148].

Initially, damaged melanocytes produce some slightly mutated but normal looking cells. At this stage, the lesion is very difficult to detect. Examining under a microscope, one may detect a group of normal looking cells with excessive mitosis (cell division). Over a period of time, if the genetic disorder increases, some new cells show abnormal shapes and orientations. Clinically, the lesion is called an atypical nevus. The increasing genetic disorder is reflected by colour variation and irregular shape. The lesion may eventually be larger than 6 mm in diameter. Even though atypical nevi are potential precursors and markers for melanomas, many of them remain clinically stable and never progress to melanomas [14]. However, with more gene mutations, some atypical nevi evolve continuously to melanomas with more abnormal cells. At this stage, the lesion is called in-situ cancer, where the lesion is entirely in the epidermis. The clinical distinction between atypical nevi and in-situ melanoma can be subtle. Melanoma in-situ is a flat and thin lesion. The lesion may enter a period of horizontal growth phase [12, 99]. Due to the variable growth rate of the atypical melanocytes, the in-situ lesion may exhibit a markedly asymmetrical shape with indentations and protrusions. A protrusion usually indicates the rapid growth of a sub-population of the melanocytes. Furthermore, the uneven distribution of melanin may cause multiple shades of brown or red-brown colour [14]. Melanoma in-situ is a non-invasive early stage cancer, which may remain contained indefinitely.

Through blood or lymphatic vessels, malignant cells are capable of establishing a colony in other body parts. Once having broken away from the basement membrane of the epidermis, the cancerous lesion enters an aggressive vertical growth phase and quickly invades the papillary dermis, reticular dermis and subcutis [12, 99]. The malignant cells can also spread to downstream lymph nodes or blood vessels to disturb the vital function of distant internal organs, such as lungs [127]. Sometimes when the body's immune system responds to a new cancer by

attacking and destroying the abnormal melanocytes in the dermis, the lesion shows prominent notches and indentations, called regression [12, 99, 101, 124].

Many staging systems have been developed to classify the disease and to guide disease management [12, 14, 124]. The best prognosis factor for a localized primary melanoma is the thickness of the lesion³, which is inversely proportional to the survival rate. (See Table 1.1 for the 10-year survival rate published by B.C. Cancer Agency (BCCA).) A primary lesion with thickness less than 1.0 mm is usually considered as a minimal or a low risk lesion due to a low chance of metastasis. However, the cure rate drops significantly when the lesion thickness is over 1.0 mm. The main treatment for all stages of melanoma is surgical excision with an increasing margin for increasing lesion thickness. Other treatments for possible metastases include regional lymph node dissection or radiation, and chemotherapy [12, 14, 124].

Table 1.1. Melanoma thickness and 10 year survival rates (BC Cancer Agency 1999) [14]

Lesion thickness	Survival rate (%)
< 1.0 mm	90
1.0 - 1.49 mm	78
1.50 - 4.00 mm	65
> 4.0 mm	40

1.3. Public health programs

The attempt to diagnose melanoma as early as possible has led to public education and screening programs in many countries [80, 97, 115, 125]. Research programs promoting self-screening for the high-risk groups are also underway [15]. Two well-known self-screening guidelines are the ABCD rule [44, 57] and seven-point checklist [57, 62]. The ABCD rule, promoted by the America Cancer Society, describes the clinical features of melanomas using

³ The thickness of a lesion is measured from the top of the skin, the granular layer of the epidermis, to the base of the lesion by an ocular micrometer.

mnemonics: A (asymmetry), B (border irregularity), C (colour variegation) and D (diameter \geq 6mm). On the other hand, the seven-point checklist, advocated by a group of dermatologists from Glasgow, emphasizes the progression of the symptoms. This checklist consists of three major features (change in size, shape and colour) and four minor features (inflammation, crusting or bleeding, sensory change, and diameter \geq 7mm). When any of the major features is detected in a melanocytic lesion, immediate help from health professionals is recommended. The presence of any minor features is advised to be monitored regularly.

The success of public campaigns for melanomas has also resulted in a flood of patients with thin lesions to general practitioners. Most of the general practitioners are not well equipped to deal with these very thin lesions, which are difficult to diagnose clinically. Therefore, to improve diagnostic accuracy and to relieve the workload of physicians, research studies for non-invasive diagnostic devices have been carried out in many directions, such as lesion progression tracking [18, 60, 136, 145], high-frequency ultrasound detection methods [36, 50, 71], and computer-aided diagnosis [21, 27, 40, 58, 59, 73, 138-140]. In particular, computer-aided diagnosis shows potential to differentiate melanomas from nevi based on their characteristics.

1.4. Non-invasive computer-aided diagnostic systems

A non-invasive computer-aided diagnostic system typically consists of several components: image acquisition, image processing, and a classifier with a knowledge database. (See Fig. 1.4.) When a melanocytic lesion is captured in vivo as a digital image by either scanning a colour slide or using a digital camera, the characteristics of the lesion, represented by a set of numbers called feature scores, can be extracted from the digital image by the image processing component. Feeding the features to a classifier which is connected to a medical knowledge database can generate a computerized diagnosis, suggesting whether the lesion is benign or malignant.




Figure 1.4. Schematic diagram for a non-invasive computer-aided diagnostic system

1.4.1. Image acquisition

Normally, the subsurface structures at the dermoepidermal junction and the superficial dermis cannot be examined by the naked eye or by microscopes because of the surface reflection (also called regular reflectance) at the air/skin interface [1, 3, 74]. Therefore, imaging melanocytic lesions with a camera or a microscope results in surface view images. These images show all four of the clinical features for the lesion, which can be described by the ABCD rule, namely asymmetry, border irregularity, colour variegation and diameter.

To examine the subsurface features *in vivo*, a special procedure or some special equipment is required. There are two common techniques to capture the subsurface view.⁴ The first technique is known as epiluminescence microscopy (ELM), which applies a drop of oil on the lesion before it is examined under a dermatoscope or a binocular stereomicroscope⁵ [11, 16, 74, 116]. The oil makes the epidermis transparent by reducing the refractive-index

⁴ Researchers have been experimenting with different *in vivo* techniques to extract subsurface features. Encouraging results have been reported recently by using multi-band spectrophotometric intracutaneous analysis [106].

⁵ The technique is also called dermoscopy, dermatoscopy, or surface microscopy.

mismatch between the corneum layer of the epidermis and air [74]. Then a glass slide is placed on top of the oil and the subsurface features of the lesion are examined under a magnifying device with magnification ranging from 10X to 40X. However, air bubbles are often trapped underneath the glass slide which may interfere with automatic analysis procedures. A new technique to capture subsurface features employs two polarized filters⁶ [67, 130, 131]: one of them is attached to the light source and the other is attached to the viewing device. (See Fig. 1.5.) The filter at the light source ensures the incident light has the same phase angle. When the polarized incident light reaches the skin, part of the light is reflected by the air/skin interface and the remaining part enters the skin. The reflected portion is called regular reflectance which has the same phase angle as the incident light. The penetrated portion is further absorbed and scattered by the epidermis and dermis, and becomes diffuse with a complete random phase angle. Finally, the diffused light is reflected out of the skin by white collagen fibres at the dermis [1, 34]. Before the re-emitted light of the regular reflection and the diffused penetrated portion reaches the viewing device, the re-emitted light is filtered by the second polarized filter, which ensures that only the re-emitted light with a certain phase angle is detected. When the viewing device filter is set to the same phase angle as the source filter, a surface view is obtained by capturing the regular reflection⁷. A subsurface view is obtained by setting the viewing device filter perpendicular to the source filter to block out the regular reflection entirely.

Both oil/glass slide and polarized filters techniques produce a similar and comparable images [130, 131]. In the subsurface view, not only the ABCD features can be clearly seen⁸, but the subsurface structures at the dermoepidermal junction and the superficial dermis are also visible. The most important subsurface structures are the pigment network, brown globules,

⁶ The technique is sometimes called digital videomicroscopy.

⁷ Part of the reflected light from the diffused, penetrated portion is also captured. However, surface view features are clearly visible when the two polarized filters are set to have the same phase angle.

⁸ The geometry properties (asymmetry, border shape and size) and colour property of the lesions are presented in the subsurface view, even though lesions manifest different colours in the surface view and the subsurface view.

black dots, radial streaming and pseudopods [6, 11, 74, 116]. (See Fig. 1.6.) The pigment network is a brown colour line network over a diffused brown colour background of the lesion. These brown colour lines are caused by the melanin deposited along the vertical edges of the convoluted papillary dermis, called rete ridges [120]. Brown globules are the melanocytic nests around the papillary dermis, and black dots are chunks of melanin in the epidermis. Radial streaming and pseudopods are the brown or black line patterns and dots, respectively, at the margin of the lesion. These subsurface features can be used to augment the ABCD features in diagnosing melanomas [17, 40, 41]. An irregular pigment network or brown globules, and the existence of radial streaming and pseudopods are associated with melanomas. However, interpreting subsurface structures is difficult and subjective. Only dermatologists properly trained with the subsurface view recognition technique can benefit from subsurface features to improve the diagnostic accuracy. Untrained dermatologists tend to have difficulty in distinguishing subsurface features, resulting in a low diagnostic accuracy. [16, 122].

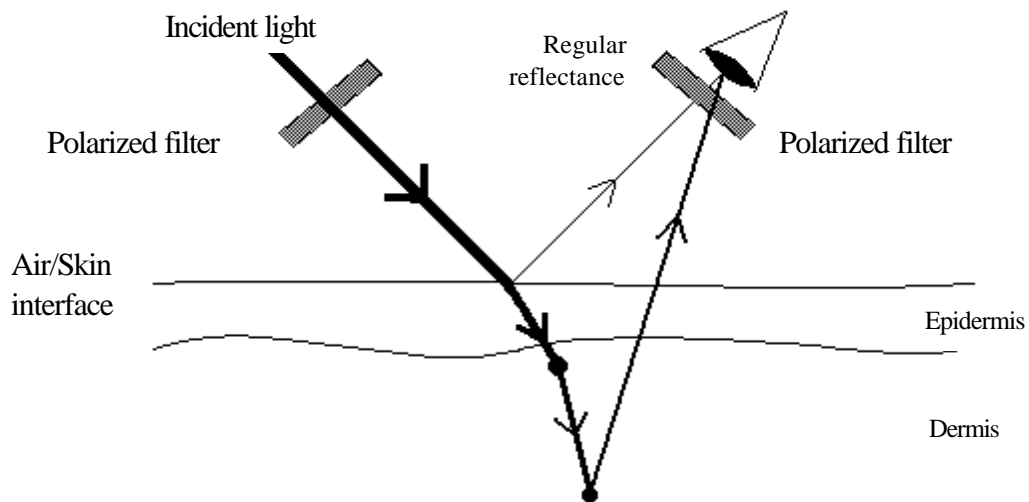


Figure 1.5. Obtaining subsurface view using two polarized filters



Figure 1.6. Subsurface view of a melanocytic lesion. The pigment network is clearly seen.

1.4.2. Image processing

Most experimental computer-aided diagnostic systems process and extract surface view (ABCD) features from skin lesion images, which can be obtained from either the surface view examination [4, 129, 149] or the subsurface view examination [130, 131]. Fully automatic image processing systems on surface view features have been reported [56, 81]. On the other hand, there has been very little work done on automatic subsurface feature extraction at the present time.⁹ Subsurface features are often evaluated subjectively and entered manually in to a classifier [17].

1.4.3. Lesion classification

Various classification methods based on well-developed theories have been examined. The popular neural networks [17] simulate the human neural system. These neural networks are usually treated as *blackboxes* and their classification rules cannot be interpreted with the input features. Other classification techniques have also been investigated, such as AI-based expert systems [35, 77, 133], statistical based methods [56, 129, 130] and principal component

⁹ Automatic detection of pigment networks, brown globules and radial streaming have been reported recently [25, 40, 41, 132]. More algorithms are expected in the future.

transformation methods [33]. Theoretically, classification results from these methods can be interpreted in terms of the input features, but often the interpretation is extremely difficult, especially when the input feature set is large.

1.5. Border shape as a defining feature

The computer-aided classification paradigm is promising when a large amount of medical knowledge is accumulated. However, one of the crucial factors for a system to succeed is to select a set of features that provide good separation between nevi and melanomas. According to the two best-known clinical guidelines (ABCD rule and the seven-point checklist), the lesion border *shape* is one of the important features. In particular, irregularities (indentations and protrusions) along a lesion border often reflect the genetic instability of the lesion.

When a lesion border is studied carefully, we notice that there are two types of irregularities: texture and structure irregularities [31]. Texture irregularities are the fine variations along the lesion border, while structure irregularities are general undulations of the perimeter. Fig. 1.7 shows examples of both types of irregularities. Lesion A has no structure protrusion and indentation, but a lot of texture irregularities. Lesion B shows a structure protrusion at the top of the border but has less texture irregularities than the other two borders, while lesion C has a prominent structure protrusion and indentation at the bottom of the border. Detecting and measuring texture irregularities may be subject to noise of the hardware imaging devices and/or software programs. On the other hand, structure irregularities may infer unstable melanocyte growth pattern or regression of a melanoma, and have been reported to have a higher correlation with melanomas than texture irregularities [31]. Therefore, the success of the computer-aided diagnostic system depends upon the ability to thoroughly assess the border shape of melanomas, especially, the structure irregularity.



(a)



(b)




(c)



(d)



(e)



(f)

Figure 1.7. Three lesions and their borders. (a) and (b) Lesion A and with its border outlined. (c) and (d) Lesion B and with its border outlined. (e) and (f) Lesion C and with its border outlined. The procedure for extracting lesion border is discussed in Chapter 3

1.6. Outline of the thesis

Chapter one states the motivation of the thesis: measuring the lesion border shape accurately may improve the effectiveness of a computer-aided diagnostic system for malignant melanoma.

Chapter two reviews the previous work on measuring border irregularity of melanocytic lesions. Because shape analysis has been an active research topic in computer vision, methods related to the development of the new shape measure are also examined.

Chapter three presents the data collection and the pre-processing steps for the new shape measurement algorithm. A full spectrum of melanocytic lesion images, ranging from clinical benign lesions to malignant melanomas, has been collected. Before these lesion images can be used for border shape analysis, they have to be processed by two programs. First, the dark thick hairs of the skin images, which may interfere the subsequent automatic segmentation program, are removed by the software program called DullRazor. Then the lesion border is extracted automatically from the skin image.

Chapter four describes the methodology of the new border irregularity algorithm, which measures all indentations and protrusions along the lesion border in a multi-scale environment. Two area-based indices generated from individual indentations and protrusions are important measures for border irregularity that can be used as input features for a computer-aided diagnostic system.

Chapter five reports a user study for validating the two new measures discussed in chapter 4. The user study result showed that the new measures vastly outperform other border shape descriptors and may be useful for diagnosing the malignancy of a lesion.

Chapter six concludes the thesis with discussions on the contribution of the research work and the future research directions.

Chapter 2

Related Work

2.1. Published border shape descriptors for melanocytic lesions

Because the irregularity of a lesion border has a strong correlation with the genetic stability of the melanocytes, all computer-aided diagnostic systems attempt to measure the border shape. However, there is no *standard* algorithm for such a measurement, although several methods have been investigated.

2.1.1. Global measures based on a single-scale

One of the most popular measurements is called compactness index (CI), which estimates the *roundness* of a 2D object [33, 38, 51, 140, 149]. Since circles are the most compact 2D objects, they have the smallest index of 1. The index for other 2D shapes is greater than 1 and is computed by the following equation:

$$P^2 / 4pA, \quad (2.1)$$

where P is the perimeter of the object and A is the object area. The popularity of CI is due to its computational simplicity; however, there are two drawbacks for this measure. It is very sensitive to noise along the border. In particular, the square term of the perimeter used in Equation 2.1 amplifies the noise. More importantly, CI cannot detect structure irregularities. Object with different shapes can associate with the same index.

Guthowicz-Krusin *et al.* [56] exploited another property of a 2D object to infer its border shape. For a lesion with irregular shape, there is a large variance in the radial distance, the distance between its centroid and border. Hence, Guthowicz-Krusin *et al.* estimated border irregularity by analysing the variance of the radial distance distribution. However, the measure is unstable because the centroid location is very sensitive to noise along the lesion

border. A small change to the border may move the centroid location and may completely alter the radial distance distribution.

Bono *et al.* concentrated on the length of indentations along the border by computing the ratio of the convex hull length and the perimeter of the lesion [20], but this measure does not account for protrusions of the lesion border, which is also an important clinical feature.

Andreassi *et al.* [4] attempted to use two indices, form factor and circularity factor, to capture the border shape. Form factor is a variation of CI and is defined as the ratio of lesion area over the normalized perimeter. This measure also shares the same drawback as CI; in particular, both measures cannot detect structure irregularities. Circularity factor is defined as the absolute area difference between the lesion and a circle centered at the centroid of the lesion and having the same area size as the lesion. However, circularity factor has the same drawback as Guthowicz-Krusin *et al.*'s method [56] that the computed result may be unstable due to its dependence on the centroid location.

2.1.2. Global measures based on multi-scale

Extending from single-scale methods, Seidenari *et al.* [130, 131] introduces a dual-scale index called fractality factor, which is defined as the normalized ratio of two border lengths measured by two different meters. However, implementation details such as the meter size is not reported; therefore, the algorithm cannot be reproduced and commented further.

Fractal dimension (FD) is an elegant multi-scale method that has been shown to possess a strong correlation with human's intuitive notion of roughness for curve lines [118]. The idea is based on Mandelbrot's observation that contradictory results are obtained when manmade objects and natural objects are measured using rulers of various sizes [100]. For a manmade object, the measured length L converges to its true value as the ruler size r decreases:

$$L = N(r)r, \tag{2.2}$$

where $N(r)$ is the number of rulers of size r required to cover the border. However, for natural objects such as coastlines, which possess self-similarity¹⁰ and randomness properties, the measured length L increases as the measuring unit size r decreases. The surprising inverse relationship between L and r is due to the fact that many bays and promontories smaller than the ruler size r are unnoticed and omitted; they only become noticeable if the unit size r decreases. Mandelbrot formulated the inverse relationship between $N(r)$ and r as:

$$N(r) = I r^{-D}, \quad (2.3)$$

where I is a scaling constant and D is the characteristic of the coastline known as fractal dimension. The value of D is a fractional number that is larger than or equal to the Euclidean dimension of the object. For example, a straight line and an open circle have a fractal dimension of 1. The fractal dimension increases monotonically and approaches 2 when the roughness of the curve increases. FD has been used to estimate the irregularity of a lesion border [31, 32, 57, 58, 108]. It has also been applied in analysing natural object textures [28], mammographic density patterns [26, 119] and cervical cells [96].

Even though FD is designed to measure the overall roughness (jaggedness) of a border, FD is insensitive to structure features and, hence, it is not suitable to measure structure irregularities of a lesion border. For example, the FD of an open circle is same as a straight line (which equals to 1). In other words, FD is insensitive to a smooth curve with a constant curvature. The problem is also demonstrated by computing the FD of two phantoms and two lesion borders, shown in Fig. 2.1. The phantom C1 has a larger indentation than C2, but they have the same FD values. More surprisingly, the lesion border L1 has a slightly larger FD (1.16) than L2 (1.12) even though L2 has a prominent protrusion and indentation at the bottom of the lesion. The problem arises because L1 is more jagged (has more texture irregularities) than L2. From these two phantoms and two lesion borders, it is clear that the FD measure is incapable of detecting structural features.

¹⁰ Self-similarity is an extension of the mathematical notion of similarity: two objects are similar if they have the same shape, regardless of their size [117].

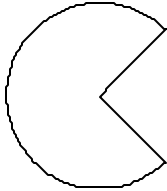
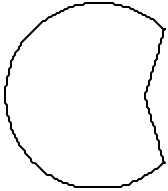
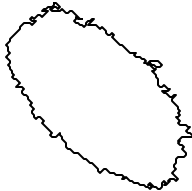

		
Phantom	C1	C2
Fractal dimension	1.02	1.02
Structure fractal dimension	1.19	1.18
		
Lesion border	L1	L2
Fractal dimension	1.16	1.12
Structure fractal dimension	1.28	1.21

Figure 2.1. Fractal dimension and structure fractal dimension of two phantoms, C1 and C2, and two lesion borders, L1 and L2.

In an attempt to capture major structural features, Claridge *et al.* [31] designed a measure called Structure Fractal Dimension (SFD). The SFD is computed based on Equation 2.3 too, but only rulers longer than a predefined threshold length are included in the computation.¹¹ The major difficulty for SFD is to properly determine the threshold ruler length for all lesions. As a result, SFD does not fully capture the structure irregularity of the lesion border as shown in Fig. 2.1.¹²

¹¹ In their implementation, Claridge *et al.* simulated the measuring process of a lesion border with dilation operations. The lesion border is covered by a series of overlapping discs with radius denoted by s . The border length is estimated by dividing the dilated border area by s . By varying the disc radius s and ensuring s is greater than a predefined threshold, the border lengths with their corresponding s values are used to determine the SFD of the lesion border.

¹² Our implementation of FD and SFD are discussed in Section 5.1.4.

None of the above methods, described in both Section 2.1.1 and 2.1.2, capture the proper notion of border irregularity for melanocytic lesions. In particular, they all have problem in detecting structure irregularities. Many methods are also sensitive to noise and may not be stable. A review of other shape algorithms in the field of computer vision and multimedia may help design a proper measure for border irregularity.

2.2. Other shape algorithms in computer vision

Analysing shape has always been an active research area for computer vision. Studying these algorithms provides some insights to the task of designing a new method for measuring the border shape of a melanocytic lesion. This section provides an extensive overview of shape algorithms in the field of computer vision which may be useful for such a task.

Shape analysis programs can be classified in many different ways according to shape attributes and analysing techniques. The most common classification scheme divides the programs according to their input types [95, 113, 137]. A program is called contour-based or external-based, when only the boundary of the object is utilized. If the interior of the object is also analyzed, the algorithm is called structural-based or internal-based. Classification schemes can also be made based on the internal shape representation of an object [113]. The scalar transform technique represents the intended object shape quantitatively by numeric scalars or vectors, called shape descriptors, which are often analyzed subsequently by statistical or neural network methods. A space-domain technique produces non-numeric and graphic representations. Shape analysis programs can also be dissected in other ways: global features vs. local features and single-scale vs. multi-scale [8, 75, 137]. Algorithms based on global features tend to be simple, but they may be unstable as a small change in the input shape may drastically alter the analysis output [22, 95]. The instability may be alleviated by introducing local information, at the expense of increasing programming complexity. Moreover, local information may be structured to facilitate the analysis. Similar comparisons can be made for single-scale and multi-scale algorithms. Multi-scale programs are harder to implement but they return more information than single-scale programs.

2.2.1. Parameterization of object borders

When an object border is analysed, a well-known technique called parameterization can be applied to reduce the dimensionality of an object from 2D to 1D. There are many ways to implement the parameterization technique. The radial function used in the study of Guthowicz-Krusin *et al.* [56] is one example. A set of radial vectors with equal angular spacing is extended from the object centroid to its boundary. The length of the radial vectors forms the radial function $r(\mathbf{q})$, where \mathbf{q} is the regular sampling angle. Although the radial function is easy to compute, the function may be unstable.

Another way to implement the parameterization technique is to define the border function $f(t)$ as two linear functions x and y :

$$f(t) = (x(t), y(t)), \quad (2.4)$$

where $0 \leq t \leq 1$ is the path length variable along the border, and $x(t)$ and $y(t)$ are the one dimensional function of $f(t)$ along the x and y coordinates [103]. When the boundary forms a closed contour, $f(t)$ is periodic. There are many variations for the formulation of Equation 2.4. Instead of using two linear functions $x(t)$ and $y(t)$, Freeman's chain code defines $f(t)$ as the direction of the pixel t with respect to its predecessor at $t-1$ [42, 43]. Zahn and Roskies define $f(t)$ as the net amount of angular bend between the initial point of the parameterization process and the point t [154]. All three variations of $f(t)$ are equivalent and they have often been used as the shape representation of an object border. In this thesis, we used the formulation of Equation 2.4 to parameterize a lesion border, since the curvature of the border can be derived directly from the equation.

2.2.2. Shape descriptors

Describing the object shape as numeric shape descriptors is more appropriate than the non-numeric, graphic representation of the space-domain technique, because the numeric descriptors can be analysed easily by most skin lesion classifiers. There have been many published methods for computing shape descriptors. The well-known descriptors are fractal

dimension (which is described in Section 2.1.2) Zahn and Roskies' Fourier Descriptors [154] and moments.

2.2.2.1. Fourier Descriptors

To compute the Fourier Descriptors, a 2D object boundary is parameterized by a cumulative function $f(t)$, which is defined as the net amount of angular bend between the point t , $0 \leq t \leq 2\pi$, and the initial point of the parameterization process, $t = 0$. Assuming the boundary forms a clockwise-oriented simple closed curve with length L , it is seen that $f(0) = 0$ and $f(L) = -2\pi$. A normalized cumulative angular function $f^*(t)$ can be defined as

$$f^*(t) = f\left(\frac{Lt}{2p}\right) + t. \quad (2.5)$$

Expanding the normalized cumulative angular function in Fourier Series using polar coordinates, Equation 2.5 becomes

$$f^*(t) = \mathbf{m}_0 + \sum_{k=1}^{\infty} A_k \cos(kt - \mathbf{q}_k), \quad (2.6)$$

where the coefficients A_k and θ_k are known as the k^{th} harmonic amplitude and the phase angle of the function $f^*(t)$, respectively. Together they form the Fourier Descriptors of the object boundary.¹³ Because the harmonic amplitudes A_k 's are invariant under translation, scaling, and rotation, they have been used to describe the object shape.¹⁴ The lower orders of the amplitudes express the overall structural shape of the closed curve, while the higher orders convey information about the fine details of the curve. However, there is no natural cutoff point for the high and low order harmonic amplitudes, and small changes along the curve may produce a completely different set of Fourier Descriptors. Another disadvantage is that spatial information cannot be retrieved in the Fourier domain.

¹³ Fourier Descriptors can also be computed from other functions such as angular function, curvature function, or radial function.

¹⁴ The phase angle θ_k also conveys the shape information about the closed curve, but it is not invariant under rotation transformations.

2.2.2.2. Moments analysis

Moments and their derivatives have been widely used for object-based recognition and reconstruction. The analysis is often carried out by the object's internal pixels [143], but the boundary pixels can also be used [146]. Mathematically, the raw (p,q) th order moment of a 2D image $f(x,y)$ is defined as

$$m_{pq} = \iint dx dy f(x,y) x^p y^q. \quad (2.7)$$

where $f(x,y)$ are the internal or boundary pixels. In the traditional definition with internal pixels, the low-order moments have well-understood physical meanings. For example, the zero-order moment m_{00} is the total image power.

For a binary image with the object denoted by 1 and the background by 0, m_{00} is the area of the object. The first-order moments, m_{10} and m_{01} , can be used to compute the centroid location of a binary image. The x coordinate of the centroid, x_c , equals to m_{10}/m_{00} and the y coordinate, y_c , equals to m_{01}/m_{00} . The second order moments, m_{20} , m_{11} , and m_{02} , depict the size and orientation of the object. They can be used to construct the principal axes of the object. However, no simple physical meanings can be attached to higher-order moments.

Although infinite numbers of moments are required to portray an image precisely, only few moment terms are required for most recognition tasks. For example, recognizing alphabetical characters requires up to third-order moments, while identifying airplanes requires fourth-order to sixth-order terms. However, a reasonable reconstruction of a binary alphabetic character of 21 x 21 pixels requires all the moments up to 15th order terms [143].

Unfortunately, the raw moments defined by Equation 2.7 do not retain the same values when the imaged object is translated or scaled. With appropriate normalization, central moments and normalized central moments are often used to overcome these problems. Hu proposed seven moment invariants, which are the linear combination of the central moments up to the third-order to describe the object shape [29, 65]. These 7 shape descriptors are proven to be invariant to scaling, translation and rotation. The complex Zernike moments provide a

general solution for rotation invariant property [143]. In the language of linear algebra, Equation 2.7 is the projection of the image function $f(x,y)$ onto the basis set $\{x^p, y^q\}$, which are not orthogonal. Applying an orthogonal basis set such as complex Zernike polynomials, a set of rotation invariant moments can be obtained. The Zernike moments respond to a rotation with a phase shift and a constant magnitude, which can be used as an invariant shape descriptor.

Many other shape descriptors can be derived from moments. Two of the well-known ones are eccentricity and inertia shape [96, 137]. Eccentricity estimates the ratio of the major axis and the minor axis of the best-fit ellipse over the object. The axes are determined by the second-order central moments of the binary image function $f(x,y)$. Circles give the minimal value of 1. On the other hand, inertia shape estimates the roundness of an object by a normalized second-order moment. For a binary image function $f(x,y)$, inertia shape is defined as:

$$2p \iint dx dy f(x, y) d_c^2(x, y) / m_{00}^2, \quad (2.8)$$

where d_c denotes the distance of an image pixel to the object centroid and the first-order raw moment m_{00} is a normalization factor. A circle again has the minimal inertia shape value of 1.

Similar to many border irregularity descriptors and Fourier Descriptors, the disadvantage for moments and their derivatives is that they are constructed from global shape features under a single-scale environment. A small change in shape may lead to a completely different set of moments. This disadvantage can be minimized if the object is divided into parts. Effects of a small shape change is limited to a part of the object and stability of the measure can be achieved. Furthermore, object partitioning may form some structured organization which can facilitate the analysis.

2.2.3. Object partitioning

2.2.3.1. Skeleton based techniques

Partitioning an object into parts and deriving shape information from the partitioned parts can provide much needed local information for shape analysis programs. Blum and Nagel [19] proposed a grassfire model, called symmetric axis transform (SAT), to divide and to represent an 2D object. The SAT can be explained by a grassfire analogue. Imagine the boundary of a piece of dry grassland surrounded by non-flammable wet grassland to be ignited simultaneously. The fire propagates towards the centre of the grass block at a constant speed. The meeting point of two or more fire fronts from different directions is called a quench point and it has an equal distance to each originated boundary of the fire fronts. The quench point can also be interpreted as the centre of a maximal inscribed disc that touches two or more boundaries. The union of all such maximal inscribed discs represents the object. Furthermore, the loci of all quench points form the symmetric axes.¹⁵ Symmetric axes and their corresponding radial functions are called SATs. When points on the maximal inscribed discs other than the centres are traced, the loci form variations of symmetric axes such as smoothed local symmetries [23] and process inferring symmetry analysis [91, 92].

Symmetric axes can be used to segment the object at the branch points so that each region is equally divided. The segmented regions can be analyzed by the geometric properties of the axes and the boundary. Blum and Nagel categorized a region into four shape classes: worm, wedge, cup and flare. They also divided a boundary into seven curvature prototypes: left and right spiral in, left and right circular, left and right spiral out and straight.

After SAT was proposed, many researchers developed algorithms to compute the axes. For 2D discrete images, mathematical morphology is a popular technique which treats an input image object as a set of points [61]. With an appropriate disc structure element, an erosion

¹⁵ These graph-liked lines are also known as medial axes or skeletons.

operation removes the boundary pixels. Applying erosion operations successively strips off the image object in layers and finally erases the entire object. Therefore, the radial distance to the object boundary for an image pixel can be computed as the number of erosion operations required to remove the pixel. The symmetric axes are formed by pixels whose radial distances are greater than their neighbours. (See Fig. 2.2.) However, because of the intrinsic nature of discrete space, radial distances cannot be determined precisely and rotation invariance cannot be achieved. Moreover, the resultant symmetric axes are sensitive to noise along the boundary and may not be thin connected lines [82].




Figure 2.2. Computing the SAT by mathematical morphology method (a) The structure element K centered at the middle pixel. (b) The binary image A represents a simple object. (c) The radial distance of each pixel in the image A . The radial distance is computed as the number of erosion operations required to remove the pixel using the structure element K . (d) The final symmetric axes of image A . The symmetric axes consist of the points whose radial distances are greater than their neighbouring points (defined by K). Note that the resultant axes are not connected.

Many methods, including active snake contours rolling down a potential field [72, 90] and Voronoi skeletons [110], have been investigated to alleviate the problems of generating incorrect radial distances and non-thin and disconnected axis lines. However, these methods still suffer from the problem that a tiny protrusion always generates an axis, while indentations

could generate no axes. Therefore, the SAT is not suitable for analysing skin lesion borders, where an indentation may be a major feature and the lesion border is often jagged and rough.

2.2.3.2. Boundary based techniques

Another common partition scheme divides the object along the boundary. The divided segments can be analysed further or approximated by various type of polynomials such as straight lines, splines and B-splines [95, 113, 114, 137]. One of the boundary partition methods is called split-and-merge, which recursively splits a curve segment into two equal halves whenever the height of the curve segment exceeds a predefined threshold. Any two neighbouring segments are merged when their heights are within some maximum distance criterion. A variation of the split-and-merge method is to split the curve at the point that has the maximum height. However, the initial point selection is important for all split-and-merge methods as different initial points generate different partitions.

Partitioning objects based on human perception theory has received a lot of interest in computer vision. Researches from the fields of human perception, cognition and psychology have shown that high curvature points possess high information content [9]. Hoffman and Richards proposed to partition an object boundary at the concavities [63, 123]. When two 3D objects interpenetrate at random, according to the principal of transversality, the objects always meet in the contour of concave discontinuities. This heuristic works well with block object decomposition. However, for a smooth surface without discontinuity, Hoffman and Richards suggested a generalized approach. For any point p in a 3D surface, the surface normal of p can be defined as a vector pointing perpendicularly inwards to the surface. Furthermore, there is always a direction that the surface curves the most, and in the orthogonal direction the surface curves the least. These two directions are called the principal directions and their corresponding curvatures are the principal curvatures. Then the surface can be divided along the loci of the negative minima, the concavities, of the principal curvatures. When a 3D object is projected into a 2D space, negative minima of the principal curvatures are mapped to concave cusps in a 2D silhouette. Hence the 3D partitioning rule can be modified for 2D planar curves by dividing

an object boundary along the minima of negative curvature, if the boundary is traced counterclockwise.¹⁶

Siddiqi and Kimia [134] argued that proper object partition must involve both object boundary and interior. Therefore, they extended Hoffman and Richard's partition scheme and introduced the concept of part-lines, which are curved lines located completely within the interior of the object and terminated on the concavities of the boundary. Specifically, there are two types of part-lines: limbs and necks. A limb is a part-line for a fixed attachment to an object, such as bird's beak. (See Fig. 2.3a.) The tangents along a limb vary smoothly. On the other hand, a neck part-line separates an articulate part, such as a human's neck or a fish tail from its object body. (See Fig. 2.3b.) The length of a neck is the minimal diameter of an inscribed circle. To obtain all the limb-based and neck-based parts, the concavities of the boundary are first located. Then each pair of concavities must be considered separately to determine whether there is a legitimate limb or neck part-line. The notation of salience, a likelihood of a part-line, is introduced to resolve the conflict with multiple part-lines from the same concavity point. However, both Hoffman and Richards' partition scheme and Siddiqi and Kimia's part-line approach handle protrusions only. They do not address indentations along the border, which are important features to diagnosing melanocytic lesions.

Focusing on how to infer the deformation history of a 2D planar object, Leyton's process-grammar analysed every protrusion and indentation along the border [91, 92]. He conjectured that all deformable objects begin from a circular shape. Applying continuation and bifurcation forces on the border distorts the object shape into protrusions and indentations. The direction of the deformation force lies along the symmetric axis which bisects the protrusion or indentation. Therefore, by locating all protrusions and indentations and tracking their symmetric axes, one may discover the history of the deformation process. As the result, although Leyton partition all protrusions and indentations along the border, his process-grammar does not return

¹⁶ Reversing the tracing direction flips the foreground and background and reverses the sign curvature values.




Figure 2.3. Limb-based and neck-based part-lines (a) A limb-based part. The limb part-line is denoted by the dotted curve line. (b) A neck-based part. The neck part-line is denoted by the dotted line, which is the minimal diameter of an inscribed circle.

rigid curve segments nor interior parts, but rather dynamic processes, which can be used to document deformation history. Furthermore, both Leyton's process-grammar and Hoffman and Richards' partition rule consider the theoretical aspect of partitioning an object boundary. They do not address how to handle a ragged border under a multi-scale environment.

2.2.4. Multi-scale methods

Some single-scale shape analysis methods can be extended to multi-scale by including two or more resolutions. Fischler and Wolf used two measuring sticks with different lengths, two scales, to determine critical break points for curve partitioning in discrete space [39]. These break points, termed salient points, correspond to the points with high curvature values. The bi-scale algorithm is repeated for each measuring stick by sliding it along the curve pixel by pixel. The maximum deviation d from the curve to the stick is recorded. The local maxima of all the recorded d 's are the salient points for that measuring stick. The final salient points of the curve are comprised of the findings from both sticks. Similarly, many skeleton techniques can also be extended into a multi-scale environment. For example, the erosion operation can take disc structure elements of various sizes to simulate different scales.

When more than one scale is applied, the natural questions to ask are: What are the optimal scales and how to incorporate results from various scales? An object may manifest different shapes at various resolutions, and shape features may disappear at a coarse scale, or the salient locations may shift from scale to scale. Hence, selecting appropriate scales,

organizing information computed from different scales and relating information from one scale to another become important issues for multiple-scale analysis. To address the above questions, many multi-scale shape analysis methods take a systemic approach by working with all scales. Fractal dimension is one of such methods where the statistical dimension, by definition, is computed by measurements from all scales. Other well-known multi-scale methods include wavelet transform and scale-space filtering.

2.2.4.1. Wavelet transform

Another popular multi-scale method in active research is the wavelet transform, which linearly decomposes a signal $f(x)$ into multi-scales based on mother wavelet basis functions ψ_k :

$$f(x) = \sum_k c_k \mathcal{Y}_k(x), \quad (2.9)$$

where c_k are the coefficients for the transformation [2, 52, 69, 141, 142]. When sinusoidal functions are used as the basis functions, wavelet transform becomes the well-known Fourier Transform, which do not retain spatial information. However, the balance of retaining information on both spatial and frequency domains can be achieved for other basis functions with local supports, such as Daubechies, Coiflet, Harr, and Symmlet wavelets [52, 141]. These wavelets have some straight mathematical properties such that their dilations (compressions) and translations are orthogonal. A fast pyramidal algorithm has been designed for the discrete wavelet transform. The discrete signals $f(x)$ of length 2^D are convolved with a low-pass filter and a high-pass filter which are derived from the mother wavelet functions.¹⁷ The filter outputs are down-sampled and concatenated together to form the base level of the resultant pyramid structure. The high-passed outputs are not involved in further computation; however, the low-passed output is processed recursively for higher pyramid scale levels. For pyramid scale level $i+1$, the low-passed output at level i is filtered again by the low-pass and the high-pass filters, down-sampled and then concatenated. The final pyramid has D scale levels. Because of the

¹⁷ The signals f can be 1D or 2D. For a 2D discrete function, both dimensions must have size equal to power of 2.

down-sampling step, the transformed signals maintain the same size as the original signals. However, similar to other pyramidal schemes, relating information from various wavelet scales can become a complicated task.

2.2.4.2. Scale-space filtering

Witkin's scale-space theory provides another way of handling multi-scale problems that emphasizes organizing and relating information among scales and the causality property of filtering [151]. In particular, he introduces a new image space with a continuous scale parameter σ , called scale-space, to record the smoothing process of the signals by Gaussian kernels because Gaussian kernels are the only kernel possessing the well-behaved causality property: features manifest in the coarse scale must persist through to the zero-order scale [10].¹⁸ During a continuous Gaussian smoothing process, no new feature is created. For a 1D object function $f(t)$, the scale-space $F(t, \sigma)$ is defined by

$$F(t, \mathbf{s}) = f(t) \otimes g(t, \mathbf{s}) = \int_{-\infty}^{\infty} f(u) \frac{1}{\sqrt{2\pi}} e^{-\frac{(t-u)^2}{2s^2}} du, \quad (2.10)$$

where \otimes denotes the convolution with respect to t and σ is the Gaussian standard deviation that controls the amount of smoothing. Unlike the wavelet transformation, there is no sub-sampling compression between scales. Instead, information among scales is tracked by a 2D binary scale-space image, which is constructed by placing σ in the y-axis and the location of the inflection points, the zero crossing of the second derivatives of $F(t, \sigma)$, in the x-axis. (See Fig. 2.4.) These inflection points are used to represent the function shape. As σ increases, the smaller extrema of the function $f(t)$ disappear first and $F(t, \sigma)$ evolves to a smoother function where the number of inflection points decreases monotonically. The causality property of Gaussian smoothing guarantees that no new inflection points are created. Therefore, a more

¹⁸ Chen and Yen showed that morphological open operations do not introduce new zero-crossings as the smoothing operation moves to coarse scales [30, 66]; however, the morphological smoothing behaves differently from Gaussian smoothing [75, 76, 79].

significant structure of the function can survive a larger σ . Moreover, the localization problem of the inflection points can be solved by a coarse-to-fine tracking of their locations in the scale-space image back to the zero-order scale, the original function.




Figure 2.4. Scale-space filtering for a 1D function. (a) The 1D function $f(t)$ drawn at the bottom of the graph is smoothed by various Gaussian kernels with increasing σ . (b) The scale-space image of (a). The zero crossings of the second derivatives of the 1D functions in (a) are used to form this scale-space image.

Asada and Brady used a scale-space technique to pinpoint the significant curvature changes for a set of primitives (corner, smooth join, crank, end and bump) along a 2D object boundary [8]. This method, termed curvature primal sketch, computes the curvature function from the object boundary and smoothes it with Gaussian kernels of increasing σ . Then the local extrema of the first derivatives and the second derivatives of the smoothed curvature function are tracked in two scale-space images. These two images are used to search for the unique signature of the primitives. Once the high curvature points of the primitives are located, they can then be used as the knot points for spline approximation of the boundary. Xin et al. [152] extended the curvature primal sketch method to handle smoothed arcs that do not have significant curvature change.

Mokhtarian built a 2D object recognition system [102, 103] and a corner detector [105] using the scale space methodology. For the object recognition system, an object border

is smoothed by Gaussian convolutions with various σ . The locations of curvature zero-crossing points are tracked in a curvature scale-space image, which is used as the signature for the object. (See Fig. 2.5.) A special matching algorithm is designed to match against the scale-space image with a large database. For the corner detection system, the maxima of the absolute curvature are extracted at the coarse scale and the scale-space image is used for localization of these corners. Mokhtarian reported that the corner detector can detect all corners without any false positive findings.

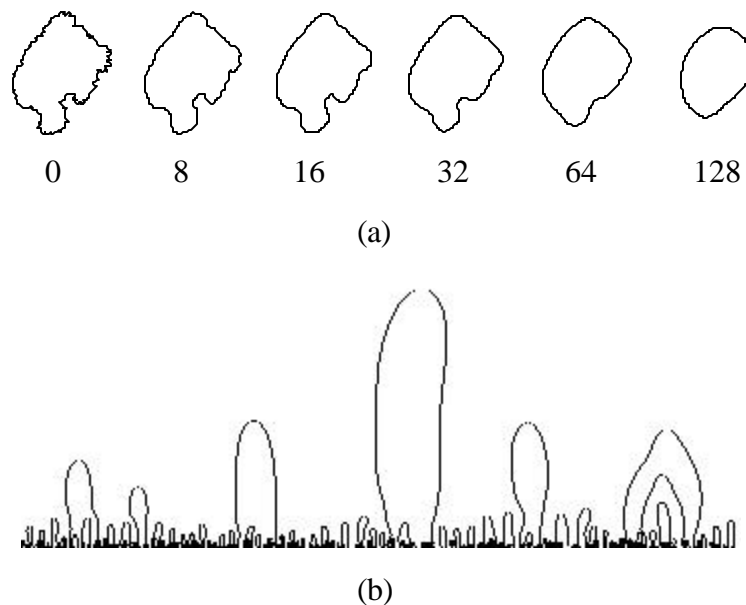


Figure 2.5. Curvature scale-space filtering for a closed contour. (a) The Gaussian smoothing can be extended to a 2D object with increasing Gaussian σ values, which are reported at the bottom of each smoothed object. (b) The curvature scale-space image of (a), with the x-axis specified by the location of the inflection points of the curvature curve.

Since no decompression is carried out between scales and an explicit data structure is used to organize information for all scales, the scale-space approach is a powerful multi-scale method to relate information between scales. Unfortunately, the existing methods are designed

to work with points along object borders, and there is no notion of protrusion and indentation segments.

2.3. Multimedia shape coding

Shape coding is an important feature in the new Moving Picture Experts Group (MPEG) standard. In MPEG-4, MPEG enhances the video compression standard of MPEG-1 and MPEG-2 by allowing the transmission of an arbitrarily shaped video object in one bit stream. The new functionality provides the infrastructure for interactive manipulation of video objects instead of pixels and paves the way for many content-based applications, such as [93].

In MPEG-4, there are two types of video objects, opaque and transparent objects. The shape of these objects is represented by a 2D binary mask, where white pixels denote the object and black pixels denote the background. Coding efficiency, scalability, error resilience and hardware complexity become important issues for a shape coder because of the potentially huge amount of video data required to be transmitted over various types of networks. MPEG-4 carefully evaluated many proposed shape coding algorithms before selecting context-based arithmetic encoder (CAE) [24] as the shape coding standard. In this section, we examine the final 4 proposed shape coding algorithms (vertex-based shape coding, baseline based shape coding, Modified Modified Reed shape coding and CAE) in the selection process.

2.3.1. Vertex-based shape coding

Vertex-based shape coding [111] codes the object contour by a set of vertices. The coding scheme is equivalent to boundary-based partitioning method discussed in Section 2.2.3.2, where the object is approximated by a polygon. The number of vertices used controls the level of lossy shape coding. Unfortunately, the optimum number of vertices is difficult to determine and their placements may alter the shape of the object, which is undesirable for analysing the object shape. For lossless shape coding, the coding scheme becomes a chain code, which has been discussed in Section 2.2.1.

2.3.2. Baseline based shape coding

Baseline based shape coding [111] is a variation of the parameterization method discussed in Section 2.2.1. First the object is rotated so that the longest axis is placed on the x-axis, which is called the baseline. The y coordinate denotes the distance between a point on the object contour and the baseline. The object contour is then sampled and the y-coordinates are differentially encoded. The contour points where the direction changes are also encoded. This coding scheme optimizes the data volume required to be transmitted over the network; however, it is inefficient to access a point on the contour. Random access for the contour points is impossible until the entire coding is decoded.

2.3.3. Modified Modified Reed shape coding

The Modified Modified Reed shape coding [153] is a block-based shape coder which enhances the standard coding method for G4 facsimile. The object shape is divided into 16 x 16 macro blocks and the changing pixels (the first white/black pixels after a run of black/white pixel block) are searched line-by-line and are encoded with respect to the reference line (the line has just been processed). Similar to the baseline based shape coding method, this method is optimized for encoding efficiency at the expense of computing efficiency. Decoding the entire macro block is required to extract the original shape expressed in the block.

2.3.4. Context-based arithmetic encoder

The CAE [24] codes all pixels of an object shape after the shape is divided into 16 x 16 macro blocks. In the intra-frame mode¹⁹, the coding is based on the local context of 10 neighbouring pixels, which account for 1024 contexts. The probability distributions of the contexts and the optimal coding are pre-computed in training sessions. The CAE method was selected as the standard shape coding for MPEG-4 because the CAE outperforms other shape coding

¹⁹ We concentrate on intra-frame coding method since there is no temporal encoding in analysing lesion border shape.

methods in coding efficiency, scalability, error resilience and hardware simplicity. However, representing an object shape by its canonical form, a binary mask, often requires further transformations to other representations such as border contours, skeletons or moments shape descriptors.

In this chapter, many object shape algorithms have been discussed; however, none of these general techniques can adequately analyse the complex features of a melanocytic lesion border. A new method is required. The new method should generate numeric shape descriptors so that they can be processed directly by a melanoma classifier. Partitioning the lesion into parts help derive important clinical information. A multi-scale method is desired to handle structural protrusions and indentations along a rugged border. Combining these general techniques, the resultant shape descriptors provide an effective measure for the border irregularity of a melanocytic lesion.

Chapter 3

Data Collection and Preprocessing

Before the new shape measure is described, the data collection and data preprocessing steps are discussed, as they are important for the development of such a measure.

3.1. Data collection

Obtaining a set of melanocytic lesion images which range from clinically benign nevi to malignant melanomas is a crucial prerequisite for developing any computerized algorithms for analysing skin lesions. Pigmented skin lesion images were collected from two sources.

3.1.1. Images from a video microscopy device

The first source was a data collection project operated by B.C. Cancer Agency and Pigmented Lesion Clinic. From 1994 to 1998, dermatologists of the clinic invited patients with clinically abnormal lesions, which required surgical removal, to participate in the project. The pigmented skin lesions specified by dermatologists were digitized by a hand-held camera (see Fig. 3.1).

The hand-held camera used in the data collection project (see Fig. 3.1a) was a video microscopy imaging device. This light-weight device was connected to a shoebox-sized main unit (see Fig. 3.1b), which was further connected to a frame grabber in a Personal Computer (see Fig. 3.1c). Inside the main unit there was a halogen bulb providing a white light source. Guided by optic fibers to the hand-held camera, the light source formed a ring within a silver-coloured, hollow cylinder attached to the front of the camera. During the imaging process, the cylinder shielded off the room ambient light. Furthermore, the cylinder was in direct contact with the patient skin to stabilize the camera against excessive lateral or vertical movement. The camera had a fixed focal length, 20 times magnifying lens so that a standard imaging environment could be achieved. Polarized filters mounted in front of the camera lens and the light source permitted subsurface view features to be examined. Because the hand-held camera is both small and light in weight, it can be moved around the patient's skin surface easily to capture

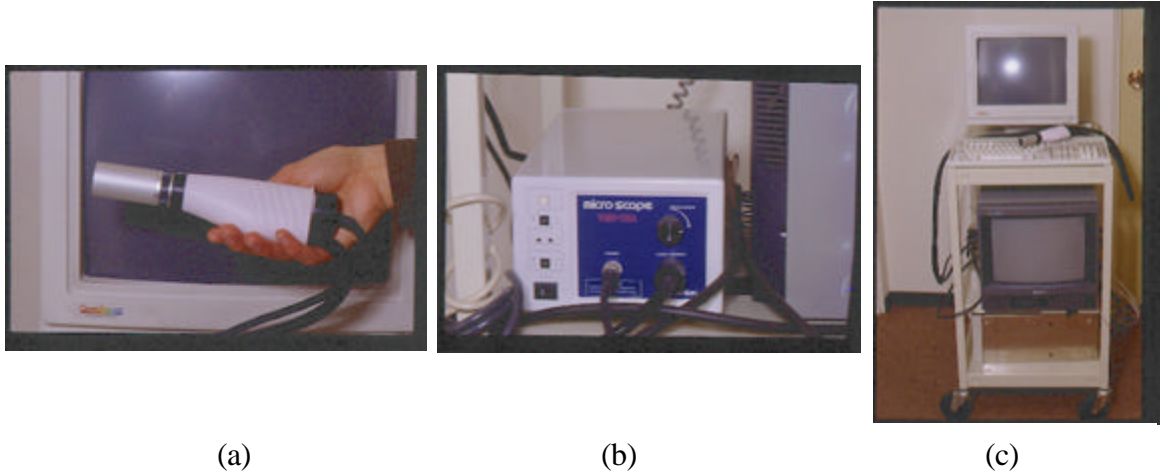


Figure 3.1. The video microscopy imaging device. (a) The hand-held video microscopy camera. (b) The main unit. (c) The Personal Computer connected by the camera.

images of skin lesions on different body parts. The resultant images were in red/green/blue (RGB) colour format, 512 x 486 pixels in size. The spatial resolution for each pixel was 0.033 mm x 0.025 mm. Every image had one lesion located near the centre and was surrounded by normal skin of differing hue. Some images also contained an additional colour marker used by the dermatologist to designate which lesion was to be imaged. The lesion could be vary in size, shape, colour and saturation. In many cases, the margin between a lesion and the surrounding skin was clinically ill-defined. Fig. 1.2a shows a skin image taken by the hand-held camera.

Follow-ups were maintained on all participants of the data collection project and the pathology findings for the surgically removed lesions were retrieved. During the project, 252 suspicious lesions from 155 patients were imaged. After the follow-up period, 178 pathology reports were collected. The histologic breakdown for the lesions included 131 pathologically benign lesions, 40 pathologically dysplastic nevi and 7 superficial spreading melanomas. We also collected 78 clinically benign lesions so that we have data for the entire spectrum of melanocytic lesions. These clinically benign lesions had no pathology report because the lesions were not removed from the patients.

3.1.2. Images from 35-mm colour slides

Another source was a set of forty-two 35-mm colour slides of malignant melanomas obtained from Pigmented Lesion Clinic of Vancouver Hospital and Health Sciences Centre. These slides were all of superficial spreading melanomas, where the entire lesion was clearly shown. Slides with non-flat body features such as digits or ears were excluded because these body parts may distort the lesion shape. The colour slides were scanned into Kodak PhotoCD digital format. Fig. 1.2b shows one of these melanoma images.

3.2. Preprocessing

Before the skin images can be used for shape analysis, they have to be processed. We use two automatic programs. First, a program called DullRazor performs dark thick hair removal [85, 86]. Then an automatic segmentation program extracts the lesion borders from the skin images [87].

3.2.1. DullRazor

Many skin images contain hairs. (Fig. 3.2 shows one of the lesion images covered by hairs.) These hairs, especially the dark thick ones with a similar colour hue to the lesion, occlude the lesion and may mislead the segmentation program. In spite of the rapid growth in the image processing applications for dermatology [40, 58, 140, 149], the hair problem has not been fully addressed. Of course, shaving the hairs before imaging sessions is one of the solutions [4, 131]. However, shaving not only adds extra costs and time to the imaging session, but also is uncomfortable and impractical especially for multiple lesions or total-body nevus imaging [136, 145]. Hence, a software approach for dark thick hair removal from skin images is needed.



Figure 3.2. A melanocytic lesion image covered by dark thick hairs

There are several approaches to handle the hair problem. Since the skin image consists of hairs, a lesion, and the surrounding normal skin, one way is to design a segmentation program to recognize all the objects. This method, however, requires a complicated segmentation program due to the fact that the hairs divide a single lesion into many sub-parts. The segmentation program must be able to join all the sub-parts together to form a single lesion. This merging process is a nontrivial task. Instead of a designing a complicated segmentation program, a preprocessor, called DullRazor, removes the thick and dark hairs from the skin images before they are segmented by a segmentation program.

DullRazor consists of three basic steps: (1) Identifying the dark hair locations, (2) replacing the hair pixels with the nearby non-hair pixels, and (3) smoothing the final result. All these steps are discussed in detail.

To locate the dark hairs, a generalized grayscale morphological closing operation is applied to the three colour bands separately [61]. The grayscale closing operation smoothes out the low intensity values, i.e., the thick dark hair pixels, along the structure element direction. The four structure elements, 0° , 45° , 90° and 135° , are shown in Fig. 3.3. The generalized grayscale closing image is obtained by taking the maximum response from the individual closing operations for each colour band. Finally, a binary hair mask image is created by thresholding

the absolute difference between the original colour band and the generalized grayscale closing image. This hair mask divides the hair and non-hair regions into disjointed areas.

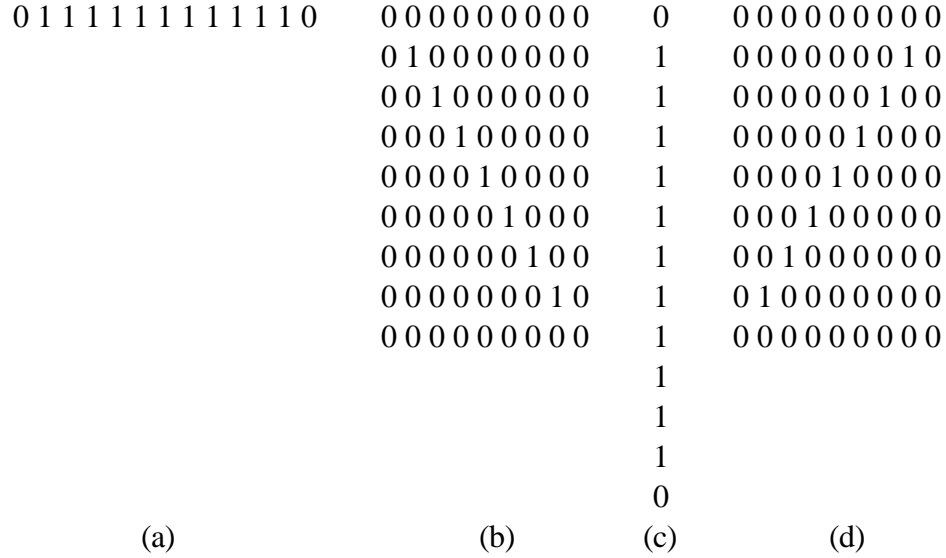


Figure 3.3. Structure element for the generalized closing operation. (a) 0° structure element, centered at (6, 0), (b) 45° structure element, centered at (4, 4), (c) 90° structure element, centered at (0, 6), (d) 135° structure element, centered at (4, 4)

Let G_r be the generalized grayscale closing image of the original red band, O_r , and S_0 , S_{45} , S_{90} and S_{135} are the structure elements in the horizontal, diagonal, and vertical directions. G_r can be expressed as:

$$G_r = |O_r - \max\{O_r \bullet S_0, O_r \bullet S_{45}, O_r \bullet S_{90}, O_r \bullet S_{135}\}|, \quad (3.1)$$

where \bullet denotes the grayscale closing operation. The binary hair mask pixel at location (x,y) , denoted as $M(x,y)$, is computed as:

$$M_r(x, y) = \begin{cases} 1, & \text{if } G_r(x, y) > T \\ 0, & \text{otherwise} \end{cases}, \quad (3.2)$$

where T is a pre-defined threshold²⁰.

A similar expression can be written for the green and blue bands and the final hair mask of the colour image, M , is the union of all three hair masks:

$$M = M_r \cup M_g \cup M_b, \quad (3.3)$$

where M_r , M_g , and M_b are the hair masks for the red, green, and blue band, respectively. Fig. 3.4 shows the hair mask M for Fig. 3.2 after the generalized grayscale morphological closing operation is applied.²¹




Figure 3.4. Hair mask M after applying an generalized closing operation to Fig. 3.2

In the second step, the binary hair mask M guides the interpolation operation to replace hair pixels by the nearby non-hair pixels. Before the replacement is performed, each pixel in the hair region of the mask M is checked to ensure that the pixel is located within a thin and long structure, i.e. the hair structure; otherwise, the pixel is rejected as noise. For each pixel inside the hair region M , line segments are drawn in 8 directions, up, down, left, right and the four diagonals, radiating from the pixel until the line segment reaches the non-hair region. These 8

²⁰ The skin images are 8 bit and the maximum intensity for each colour band, therefore, is 255. The threshold T is empirically set to 24 after testing a large set of skin images with dark thick hair covers.

²¹ Note the original mask is binary black/white, but reduction for printing makes it appears as grayscale.

line segments form 4 straight lines centered at the pixel. The length of each line is calculated and the longest one is noted. The longest line must be longer than 50 pixels and other lines must be shorter than 10 pixels. Otherwise, the pixel is rejected. The cleanup hair mask is shown in Fig. 3.5.



Figure 3.5. The cleanup hair mask of Fig. 3.4

After a pixel is verified to be inside a hair structure, the corresponding pixel in the original image is replaced by two nearby non-hair pixel values along the shortest line, the line perpendicular to the longest one, using bilinear interpolation.²² Let $I(x,y)$ be the intensity value for the replacing pixel, and $I_1(x_1,y_1)$ and $I_2(x_2,y_2)$ be the selected non-hair pixel intensities along the shortest direction. The new intensity value $I_n(x,y)$ can be expressed as:

$$I_n(x,y) = I_2(x_2,y_2) * \frac{D(I,I_1)}{D(I_1,I_2)} + I_1(x_1,y_1) * \frac{D(I,I_2)}{D(I_1,I_2)}, \quad (3.4)$$

where $D(a,b)$ is the Euclidean distance between point a and b .

Fig. 3.6 shows the resultant skin image after the replacement step. In this image, many thin lines around the edge of hair structures are still visible. The last step of DullRazor smoothes away these thin lines by an adaptive median operator. First, an extended hair mask is

²² Replacing hair pixels with spline interpolation may produce smoother results than bilinear interpolation in the expense of computational complexity. However, for the purpose of using DullRazor as a preprocessor for an automatic segmentation program, bilinear interpolation produces an adequate result.

constructed by applying a binary dilation with a 5 x 5 square structure element of all 1's centering at the middle of the square. (The enlarged hair mask for the adaptive median smoothing is shown in Fig. 3.7.) Then a 5 x 5 median filter is applied to only the enlarged hair regions, while the non-hair regions are left un-touched to preserve fine details. The final output of DullRazor is shown in Fig. 3.8.




Figure 3.6. The skin image of Fig. 3.2 after the replacement step



Figure 3.7. The enlarged hair mask for Fig. 3.5




Figure 3.8. The final output of DullRazor for Fig. 3.2

3.2.2. Segmentation

The second phase of the preprocessor task extracts the lesion border from the skin image. After occluding dark thick hairs are removed by DullRazor, skin images consist of a lesion surrounded by normal skin.²³ (See Fig. 3.8.) The intensity of the normal skin is uniformly distributed, but the lesion can vary in size, shape, colour and intensity. In many cases, the separation between a lesion and the surrounding skin is fuzzy. A three-step algorithm is used to extract the border. These steps consists of (1) a multi-stage median filter to suppress noise, (2) a process to determine the threshold value, and (3) a rule-based system to identify the lesion and to extract the border.

The first step of the segmentation algorithm removes noise caused by the imaging process while preserving image details such as fuzzy edges.²⁴ A multi-stage central weighted median filter (CWMF) [121] with a set of linear and curved filter windows is applied to the red, green and blue bands of a skin image separately. The CWMF is a special case of the well studied median filter (MF), which outputs the median pixel values in a running window of size N

²³ A few skin images also contains a colour marker used by the dermatologist to designate the lesion to be imaged.

²⁴ DullRazor focuses on hair regions of a skin image; non-hair regions are neither processed nor altered.

$x N$, where N is an odd integer [5, 68, 126]. It is known that the MF preserves edges while removing impulsive noise, but it also erases thin lines and clips sharp corners. The problematic behaviour of the MF can be alleviated by a set of linear and curved filter windows shown in Fig. 3.9. For example, the filter window $W(1)$ preserves details in the vertical direction, while $W(2)$ preserves details in the horizontal direction.




Figure 3.9. Filter windows for median filter

The CWMF puts more weight on the central pixel (the pixel being examined) by adding multiple copies of the central pixel prior to the median operation [78]. Let $X(i,j)$ be the input of the filter, $Y(i,j)$ be the output, and W be the $N \times N$ filter window with window size of $2L+1$. The output of the CWMF is denoted as:

$$Y(i, j) = \text{median}\{X(i - s, j - t), 2K \text{ copies of } X(i, j) \mid (s, t) \in W\}, \quad (3.5)$$

where K is a constant. When $K = 0$, the CWMF becomes a MF. As K increases, the CWMF puts more emphasis on preserving details until it becomes an identity filter (i.e. no filtering) when $K \geq L$.

For our automatic segmentation program, a three-stage CWMF is constructed as shown in Fig. 3.10 [121]. In the first stage, there are 8 CWMF operators. Each CWMF operator works with one of the linear and curved filter windows as defined in Fig. 3.9. The filter windows mask out the input pixels before they are sent to the operator. The output from the first stage is arranged into two groups. Combining with the original central pixel, each group is directed to a CWMF operator in the second stage. Finally, the result of the entire three-stage

CWMF operator is obtained by submitting the output of stage two with the central pixel to the single CWMF in the third stage.

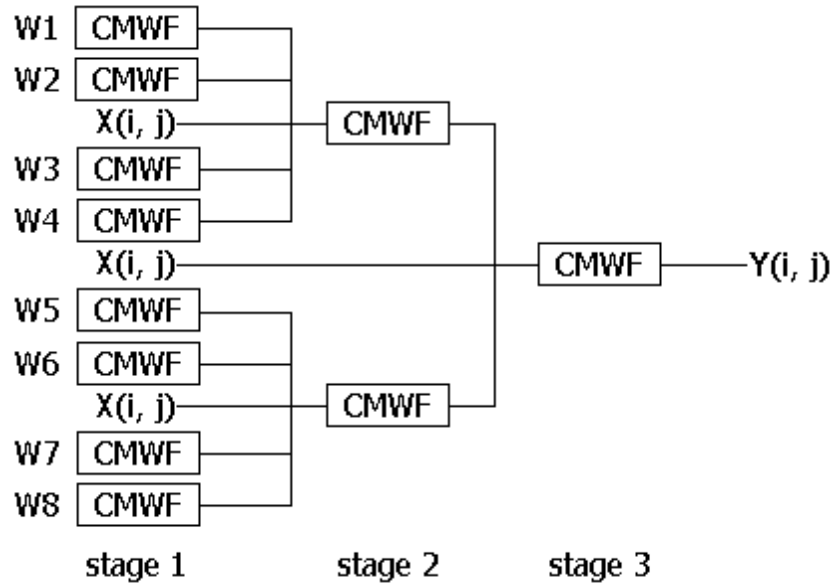


Figure 3.10. Scheme for the three-stage CWMF

For each pixel in the image, there are 11 CWMF operations, and each operation computes the median value from an extended set. An interesting property of the CWMF is that it can simplify the computation [78, 121]. Again, let $Y(i,j)$ be the output and $X(i,j)$ be the input of a CWMF operator with $2K$ copies of $X(i,j)$ being added. Also let $2L+1$ be the size of the $N \times N$ filter window, where N is an odd integer. Then the output $Y(i,j)$ can be expressed as:

$$Y(i, j) = \text{median} \{ X(i, j)[L+1-K; 2L+1], X(i, j)[L+1+K; 2L+1], X(i, j) \}, \quad (3.6)$$

where $X(i,j)[r; 2L+1]$ is the r^{th} smallest element among the $2L+1$ samples within the running window centred at $X(i,j)$.

Fig. 3.11 shows the result of applying a three-stage CWMF to the lesion image in Fig. 3.8. Because the purpose of the filtering step is to remove noise caused by the imaging process while preserving image details, minimum modification is done on the input image. The mean absolute differences are 0.13, 0.09 and 0.12, and the maximum absolute differences are 19, 20, 21 for the red, green and blue band, respectively.²⁵




Figure 3.11. Effects of the three-stage CWMF on Fig. 3.8. (a)-(c) specify the pixel locations modified by the three-stage CWMF operator for the red, green, and blue band, respectively. (d) The three-stage CWMF output for Fig. 3.8.

The second step determines the threshold values to segment a lesion from the surrounding normal skin. Three threshold values, one for each colour band, are needed and

²⁵ The maximum intensity for each colour band is 255.

each threshold value is calculated from a histogram. Since the normal skin and the lesion are the dominant objects in the image, they correspond to the two highest peaks in the histogram. Moreover, because the normal skin has a lighter colour tone than the lesions, the peak for the skin can be identified with confidence as the peak with the higher intensity value. The final threshold value T , which falls between the two peaks, is fine tuned by analysing the pixel distribution for each intensity value i between the two peaks. When intensity i is examined, all pixels corresponding to intensities $[0 .. i-1]$ are classified as base-mole, while the pixels corresponding to $[i+1 .. 255]$ are classified as base-skin. A pixel p whose intensity is i is classified as noise if there exist no connected path between p and the base-mole. In other words, there does not exist a sequence of pixels $p_0, p_1, p_2, \dots, p_n$, such that, for $j = 1, 2, \dots, n$, the intensity value of all p_j 's are within $[0..i]$, the pixels p_j and p_{j+1} are adjacent neighbouring pixels and p_n is inside base-mole. The term noise-ratio is defined as the ratio between the number of pixels categorized as noise and the total number of pixels corresponding to intensity i . The threshold T is set to the largest intensity i with an acceptable noise-ratio.

The above noise-ratio may break down if intensity i approaches the histogram peak of the normal skin, S . In this case, base-mole may occupy such a large area that there is always a connected path between base-mole and most of the corresponding pixels of intensity i . To avoid this situation, the examination of intensity i too close to S is prohibited by another ratio, skin-ratio, which is defined as the ratio between the number of pixels corresponding to i and the number of pixels classified as base-skin. With both ratios working together, the fine tuning process of T selects the highest possible i whose noise-ratio and skin-ratio are within a pre-defined range.

The final step of the segmentation algorithm combines the threshold results from the three colour bands, identifies the lesions and extracts the borders. A pigmented skin lesion is darker or redder in colour than the surrounding skin not because there is more red in the lesion, but the melanin nest absorbs most of the blue wavelength so that there is very little blue component. As illustrated by Fig. 3.12, the low intensity values in the blue band enable a stable threshold computation and so the blue band has more discernment power than the red and

green bands. For some of the very faint, low contrast images, the red and green bands may not produce a proper threshold value. This observation suggests more weight should be put on the blue band. Hence, the last step uses the blue band as the dominant band. The red and green bands attempt to capture the fine details that the blue band misses, i.e., when the blue intensity is slightly above the threshold value T and is well below S . All these considerations guide the final rule-based system to form a binary mask that divides the image into the skin region and the lesion region: the pixel at location (i,j) of the binary mask is classified as part of the lesion region if

$$B(i, j) \leq T(B) \quad \text{or} \quad (3.7)$$

$$T(B) \leq B(i, j) \leq C \quad \text{and} \quad G(i, j) \leq T(G) \quad \text{and} \quad R(i, j) \leq T(R),$$

where $X(i, j)$ is the intensity value of the band X at the pixel position (i, j) , for $X = R$ (red), G (green) or B (blue). $T(X)$ is the final threshold value of band X . C is a pre-defined parameter where $C < S$.




Figure 3.12. The red (a), green (b) and blue (b) band of Fig. 3.8

Once a binary mask of the skin lesion image is generated, the mask is passed to a blob-colouring process [13, 61] to enumerate the objects in the lesion region. Cleanup is then applied to eliminate noise objects whose sizes are smaller than some pre-defined minimum size.

The remaining objects are declared as lesions with edges outlined by a boundary tracking procedure [137]. (Fig. 3.13 demonstrates the segmented border for Fig. 3.8.²⁶)



Figure 3.13. Segmentation result. The original image is shown in figure Fig.3.2 and the image after DullRazor is shown in Figure. 3.8.

3.2.3. Issues arising in the data preprocessing phase

In the preprocessing phase, DullRazor attempts to simplify the segmentation task by removing deleterious effects of hairs. However, the software approach to hair removal has an inherent problem. It is obvious that pixel values underneath the hairs cannot be reconstructed accurately by a single view. A careful examination of Fig. 3.8 reveals traces of faded hair lines, which is why the program is called DullRazor. These traces can probably be removed, but at the cost of an excessive loss of fine details in the image. Therefore, these artifacts are not processed further prior to the segmentation program. Furthermore, other image analysis tasks, such as texture analysis or pigment network extraction that tend to be sensitive to pixel modifications, can use the hair mask as shown in Fig. 3.7 to locate the modified pixels.

Fig. 3.14 shows an example where, although most of the lesion was segmented properly, a small region at the right side of the lesion and another one at the bottom-left corner

²⁶ DullRazor and the automatic segmentation program were implemented in C using a Sun Ultra-II workstation. The execution time for the preprocessing phase was about 30 seconds.

were misclassified due to heavy hair clusters and fuzzy borders. Indeed, automatic segmentation is a non-trivial task even for skin lesion images, which have a well-defined domain of a lesion surrounded by the normal skin. Hairs and fuzzy borders increase the difficulty of the task.




Figure 3.14. A skin lesion heavily covered by dark thick hairs. (a) the original hairy image, (b) the segmentation result after DullRazor

In order to evaluate the segmentation program, Dr. David McLean and I examined all segmentation results on a colour computer monitor. First, the original skin image was shown on the monitor, and then the skin image with the segmentation result, which was superimposed as a white outline, was shown along side with the original image. By comparing these two images, the segmentation results were assessed subjectively and categorized into three groups: good, fair and poor. When the computed border outline had a good agreement with the actual lesion border, the segmentation result was classified as good. If there were some minor misinterpretations, the segmentation result was classified as fair. Finally, a poor result indicated the segmentation program missed a large portion of the actual border. After all 298 skin

images²⁷ were evaluated visually, we found that only 108 skin images, or 36%, were classified as good. Fig. 3.13 and Fig. 3.15 show two examples of good segmentation results. 144 skin images, or 48%, (see Fig. 3.14 and Fig. 3.16) were categorized as fair, and 46 skin images, or 15%, (see Fig. 3.17) were categorized as poor.




Figure 3.15. Example of good segmentation results. (a) The original image, which is also shown in Fig. 1.2a. (b) Segmentation result of (a).

Because of the complexity of a skin image, which may consist of hairs, a multi-coloured lesion, a fuzzy border and/or a low contrast image, a fully automated system that produces a perfect segmentation for all lesions is not achievable.²⁸ However, as the starting point of developing such an automatic system, the preprocessing phase produces some good results. If there is any concern with the segmentation accuracy, the extracted border can be inspected and modify manually before any subsequent analysis is applied.

²⁷ These 298 images include 78 clinically benign lesion images, 131 pathologically benign lesion images, 40 pathologically dysplastic nevi images and 7 superficial spreading melanoma images recorded by the hand-held video microscopy camera, and 42 superficial spreading melanoma images in Kodak CD format.

²⁸ A precise manual segmentation may not be achievable for a fuzzy lesion border. Increasing the resolution of the imaging device may help; however, the partial volume effect of the melanin distribution cannot be fully solved.




Figure 3.16. Example of fair segmentation results. (a) The original image (b) Segmentation result of (a). The lesion border at the bottom and at the top-right hand side is outlined incorrectly due to the subtle difference in hue between the lesion and the normal skin.




Figure 3.17. Examples of poor segmentation results. (a) A very faint, low contrast skin lesion with a fuzzy border. (b) The segmentation result of (a). (c) A lesion divided by white skin marks. (d) The segmentation result of (c).

Chapter 4

Irregularity Indices

As discussed in chapter 2.1, the previously used methods for measuring border irregularity are more sensitive to texture irregularities than structure irregularities (structural protrusions and indentations), which are the important clinical features for diagnosing melanomas. The single-scale methods compute some global properties of a skin lesion but these methods are noise-sensitive. The problem is magnified when a melanocytic lesion border is extracted by an automatic segmentation program, which tends to produce a more ragged (noisy) border than a manual segmentation program. Although fractal dimensions (FD) and structure fractal dimensions (SFD) are multi-scale methods, they cannot detect structural irregularities. To solve these problems, a stable measure sensitive to structure irregularities needs to incorporate both global and local features. Our new measure achieves this by partitioning a lesion border into protrusion and indentation segments under a multi-scale scheme. This chapter describes our algorithm in detail [84, 88].

4.1. Abstracting the lesion border

The algorithm begins with an abstraction of the lesion border, which can be extracted by a segmentation program such as the one reported in the previous chapter, as a simple closed planar curve L_0 in C^2 . Further simplification is made by parameterization of the x and y coordinates into two linear functions $x(t)$ and $y(t)$, where $0 \leq t \leq 1$ is the path length variable along the planar curve:

$$L_0 = (x(t), y(t)). \quad (4.1)$$

The initial point of the parameterization, $t = 0$, can be selected arbitrarily; however, the border is traced in the counterclockwise direction so that the interior of the lesion is in the left-hand side of the border during the tracking process. Furthermore, $L_0(x(0), y(0)) = L_0(x(1), y(1))$, because the border forms a closed contour. An example of a lesion border is shown in Fig. 4.1 with the initial point of the parameterization marked by "|".




Figure 4.1. Example of an extracted lesion border. The initial point of the parameterization, $t = 0$, is marked by "|" and the border is traced in the counterclockwise direction. The corresponding melanocytic lesion is shown in Figure 1.7e and 1.7f.

4.2. Defining indentation and protrusion segments

A simple method to measure irregularities (indentations and protrusions) along a curve is to compute the corresponding curvature function. For a point ξ in a curve $L(x(t), y(t))$, the curvature at ξ is defined as the instantaneous rate of change of the angle between the tangent at ξ and the x-axis with respect to the arc length of L . The curvature can be also expressed as $1/\rho$, where ρ is the radius of the locally best-fit circle at ξ . When the curvature is large at ξ , i.e., a small ρ , the curve turns rapidly at ξ . However, for a small curvature, i.e., a large ρ , the curve is nearly flat at ξ . It is known that the curvature function $K(t)$ can be computed as [53]:

$$K(t) = \frac{x'y'' - y'x''}{\left((x')^2 + (y')^2\right)^{3/2}}, \quad (4.2)$$

where x' and y' are the first derivatives of the functions $x(t)$ and $y(t)$ with respect to t , and x'' and y'' are the second derivatives of $x(t)$ and $y(t)$. The curvature function portrays the curve L in two ways. The sign of $K(t)$ indicates concavity or convexity of the curve at point ξ and the magnitude denotes the amount of bending. With our convention, using counterclockwise tracing

along the border and image coordinate system²⁹, positive curvature values imply concavity, while negative curvature values imply convexity. Local curvature extrema, located by the zero-crossings of the first derivative of $K(t)$, mark the tip points of concave and convex segments. These tip points are considered to have high information content [9], and have been frequently used to partition border contours into parts. Hoffman and Richards [63, 123] partitioned object borders at concave tips and their part primitives, codons, are constituted of curve segments with 0, 1 or 2 curvature extrema. Leyton [91, 92] constructed symmetric axes by bisecting concave and convex segments from their tips. He further suggested that deformation forces could act on these tip points in the direction of the symmetric axes. Siddiqi and Kimia's [134] neck-based and limb-based approach of object decomposition also put the terminals of part-lines at the concave tips.

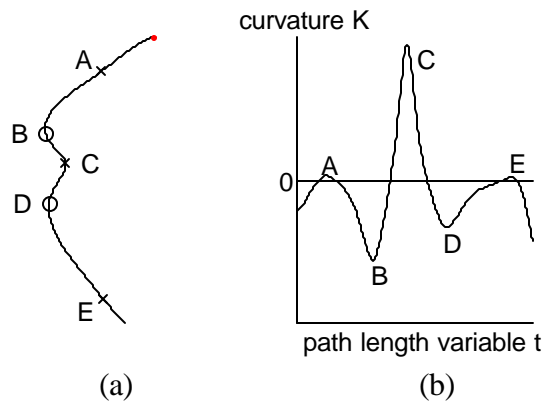


Figure 4.2. Definition for indentations and protrusions. (a) A curve segment with two protrusions and one overlapping indentation. (b) The corresponding curvature function. The points A, B, C, D, and E are the local curvature extrema.

We exploit local curvature extrema to divide the border into a set of indentation/protrusion segments (see Figure 4.2). An indentation/protrusion segment is defined as a curve segment composed of three consecutive local curvature extrema $[t_1, t_2, t_3]$, where t_1 ,

²⁹ The origin is in the top-left corner.

t_2 and t_3 are specified by the path length variables of the curve. The middle curvature extremum t_2 denotes the segment tip point and the segment type. For example, when t_2 is a concave curvature extremum, $K(t_2) > 0$, the corresponding segment is an indentation segment. Otherwise, a convex curvature extremum t_2 , $K(t_2) < 0$, specifies a protrusion segment. The local curvature extrema t_1 and t_3 delineate the extent of the segment and they have a different sign from curvature extremum t_2 :

$$\begin{aligned} \text{sign}(K(t_1)) &= \text{sign}(K(t_3)), \\ \text{sign}(K(t_1)) &\neq \text{sign}(K(t_2)). \end{aligned} \tag{4.3}$$

In other words, an indentation segment is defined as a curve segment that begins with a convex curvature extremum, following by a concave curvature extremum and a convex curvature extremum. Similarly, a protrusion segment, defined as a dual of an indentation, is a curve segment that begins with a concave curvature extremum, following by a convex curvature extremum and a concave extremum. For example, in Fig. 4.2a, a curve segment with two protrusions and one overlapping indentation is shown. The points B and D specify the tips of the protrusion segments and the point C specifies the tip of the indentation segment. The corresponding curvature function is plotted in Fig. 4.2b. The local curvature extrema, points A, B, C, D and E, are computed as the zero-crossings of the first derivative of the curvature function, K' . By applying the definition of the indentation/protrusion segments, one indentation segment [B, C, D] and the two protrusion segments [A, B, C] and [C, D, E] can be located. Therefore, a lesion border is decomposed into a set of indentation/protrusion segments by scanning the corresponding curvature function for three consecutive curvature extrema with alternating signs.

4.3. Extracting structure indentation and protrusion segments

4.3.1. Scale dependence of irregularity segments

Computing indentation/protrusion segments in discrete space using the method described in Section 4.2 is scale dependent because Equation 4.2 involves the first and second derivatives of

the functions $x(t)$ and $y(t)$, which are estimated by difference functions. For example, when $x(t)$ is transformed to a discrete function $\chi(i)$, $i = 1, 2, 3, \dots$, $x'(t)$ is approximated by:

$$x' \cong \frac{c(i-d) - c(i+d)}{2d}, i = 1, 2, 3, \dots, \quad (4.4)$$

where d is a constant. Furthermore, the parameter d controls the scaling of the difference function. When d is small, Equation 4.4 estimates x' using local information. On the other hand, global information is used to estimate x' when d is large. Therefore, estimating derivatives and, in turn, determining indentation/protrusion segments are scale dependent.

The scale dependence of Equation 4.4 raises an important issue: what is the optimal scale? One solution is to use several pre-defined scales. However, this approach has problems of selecting appropriate scales and relating information found among scales. For the new measure, a well-known method called scale-space filtering technique is used to provide a solution for scale selection and information passing in a multi-scale environment [8, 94, 105, 151].

4.3.2. Classic curvature scale-space

The idea of the classic curvature scale-space filtering [102, 103, 105] is based on continuously smoothing the original border function $L_0(x(t), y(t))$ by convolving $x(t)$ and $y(t)$ with a Gaussian kernel $g(t, \sigma)$ of increasing width:

$$\begin{aligned} L(t, \mathbf{S}) &= L_0(x(t), y(t)) \otimes g(t, \mathbf{S}) \\ &= (X(t, \mathbf{S}), Y(t, \mathbf{S})), \end{aligned} \quad (4.5)$$

where $L(t, \sigma)$ is the smoothed border, \otimes denotes an convolution operator,

$$X(t, \mathbf{S}) = x(t) \otimes g(t, \mathbf{S}), \quad Y(t, \mathbf{S}) = y(t) \otimes g(t, \mathbf{S}), \quad (4.6)$$

and

$$g(t, \mathbf{S}) = \frac{1}{\mathbf{S}\sqrt{2\mathbf{p}}} e^{-t^2/2\mathbf{S}^2}. \quad (4.7)$$

Before the smoothing process, the lesion border is parameterized (sampled) once to obtain $x(t)$ and $y(t)$. Then as expressed in Equations 4.5 to 4.7, the smoothed border is computed by convolving the original $x(t)$ and $y(t)$ with Gaussian functions of different Gaussian standard deviation σ 's, which control the amount of smoothing.³⁰ During the smoothing process, fine features are first detected when the scale parameter \mathbf{s} is small. As \mathbf{s} increases, fine features along the border are smoothed-out and the smoothed border $L(t,\sigma)$ is transformed toward an oval shape. At this stage, large global features can be extracted. The smoothing process is terminated at \mathbf{s}_{term} when all concavities of the border are removed. The smoothing process is demonstrated in Fig. 4.3 using the lesion border shown in Fig. 4.1. For this border, the smoothing process terminates at $\mathbf{s}_{term} = 129$.

The classic curvature scale-space filtering technique employs a 2D scale-space image to record certain features for the entire smoothing process in a precise format. Fig. 4.4 shows an example of a classic curvature scale-space image. The y-axis of the image represents the smoothing scale, denoted by Gaussian \mathbf{s} , while the x-axis represents the path length variable t , which specifies the spatial locations of curvature zero-crossings along the lesion border. The curvature function $K(t,\sigma)$ of the smoothed border $L(t,\sigma)$ is defined as:

$$K(t,\mathbf{s}) = \frac{\frac{\partial X}{\partial t} \frac{\partial^2 Y}{\partial t^2} - \frac{\partial^2 X}{\partial t^2} \frac{\partial Y}{\partial t}}{\left\{ \left(\frac{\partial X}{\partial t} \right)^2 + \left(\frac{\partial Y}{\partial t} \right)^2 \right\}^{3/2}}, \quad (4.8)$$

and the curvature zero-crossings are points that satisfy the following conditions along the border:

$$K(t,\mathbf{s}) = 0, \quad \frac{\partial K(t,\mathbf{s})}{\partial t} \neq 0. \quad (4.9)$$

To construct such a scale-space image, smoothed borders are analysed in turn. For each smoothed border $L(t,\sigma)$, all curvature zero-crossings are located and plotted in the 2D scale-

³⁰ The smoothed border is not re-parameterized (sampled) again.

space image. The resultant data representation is a binary image³¹ as shown in Fig. 4.4. For example, the labelled points A0 to A9 in Fig. 4.4 are the corresponding curvature zero-crossing points shown in the smoothing process of Fig. 4.3 at $s = 40$. (The curvature function is shown in Fig. 4.5.) The scale-space image captures the detected feature for the entire smoothing process.

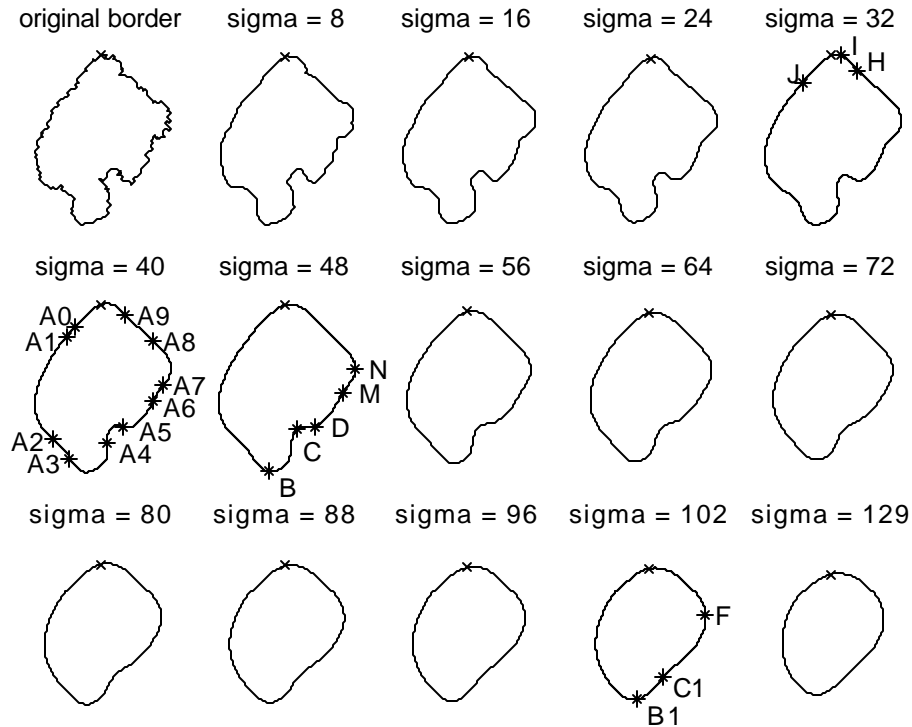


Figure 4.3. Gaussian smoothing process for lesion border shown in Fig. 4.1. Only some of the s smoothing levels, specified at the top of each subfigure, are plotted. The parameterization of the closed curves begins at the point marked as 'x' and the parameterization is done in the counterclockwise direction. The segment [H, I, J] at $s = 32$ wraps around the initial point of the parameterization process. At $s = 40$, all curvature zero-crossing points are marked as '*' and are also shown in Fig. 4.4. The protrusion segment [H, I, J] at $s = 32$, and the indentation segments [B, C, D] and [B1, C1, F] at $s = 48$ and 102, respectively, are shown in Fig. 4.8a.

³¹ The classic curvature scale-space image is considered as a binary image because it requires 2 values to encode the image. The loci of curvature zero-crossings can be coded as 1, while the background can be coded as 0.




Figure 4.4. Classic curvature scale-space image for the smoothing process shown in Fig. 4.3. At $s = 40$, the spatial locations for curvature zero-crossings are labelled as A_0, \dots, A_9 . Their corresponding points are also marked in Fig. 4.3. The point A_c , denoted the peak of a contour arc, is discussed in Section 4.3.5.

4.3.3. The sigma-ratio shape measure

The value s_{erm} , the minimum amount of smoothing required to remove all concavities on the border L_0 , reveals the overall roughness of the entire lesion border. Such a property has been recognized and evaluated in our early work towards development of a new measure for a melanocytic lesion border. We defined a simple index called Sigma-Ratio (SR) as the ratio of s_{erm} over the border length [83].³² Since a circle has a constant convex curvature with no extrema, it has the minimum index value 0. All other lesion borders have an index greater than 0 based on the roughness of the border. For example, the sigma-ratio indices for the lesions A, B, and C shown in Fig. 1.7 are 0.29, 0.53 and 0.71, respectively. (The lesion borders are shown in Fig. 4.6 again for convenience.) These ratios correlate well with the structural

³² Normalization with the border length is required to yield a scale invariant measure.

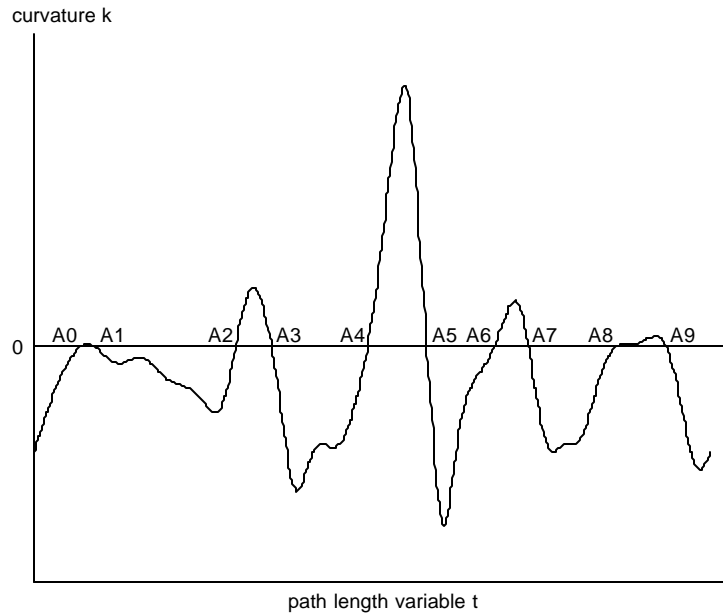


Figure 4.5. Curvature function of the smoothed border at $\sigma = 40$ of Fig 4.3. All curvature zero-crossing points are labelled as A0 to A9.

roughness of the lesion borders. Lesion border A has no structure irregularities other than texture irregularities. Lesion border B has an protrusion at the top of border, while lesion border C has major structural indentations and protrusions at the bottom of the border.




Figure 4.6. Three lesion borders of Fig. 1.7. (a) Lesion border A. (b) Lesion border B. (c) Lesion border C.

Even though this simple index works well with many lesion borders, there are some shortcomings. First the SR is non-linear. Transforming σ to $\log(\sigma)$ is required to maintain linearity [151]. A sensitivity test also shows that a very high SR value (0.96) is returned for a

circle with a long and narrow indentation such as the phantom shown in Figure 4.7a. It implies that if an occluding hair of a skin lesion is misinterpreted as a long and narrow indentation by a segmentation program, a high index value may be returned. Therefore, the SR shape index requires all hairs to be removed carefully either by shaving, or by using a software program such as DullRazor [85]. The other shortcoming is more problematic than the previous one. Since the SR is constructed from the global feature of the lesion border length, it cannot properly recognize certain structure features. In particular, the phantom with a much larger indentation shown in Fig. 4.7b has a only slightly higher SR value (0.97) than the phantom in Fig. 4.7a. These shortcomings suggest that the new measure should be area-based so that the measure is proportional to the area of the irregularity. Furthermore, the necessity of properly identifying structure protrusions and indentations is reinforced. In the next sections, we describe how to analyse the indentation/protrusion segments using the extended curvature scale-space image.




Figure 4.7. Two phantoms for sigma-ratio.

4.3.4. Extended curvature scale-space

Previously, the locations of curvature zero-crossings have been used as the detecting feature in the classic curvature scale-space filtering technique [102, 103, 105]. To analyse indentation/protrusion segments instead of border points with zero curvature values, local

curvature extrema are chosen to be the investigated feature of our extended curvature scale-space image. These curvature extrema are defined as the zero-crossings of the partial derivative of $K(t, \sigma)$ with respect to t , i.e.

$$\frac{\partial K(t, \mathbf{s})}{\partial t} = 0, \quad \frac{\partial^2 K(t, \mathbf{s})}{\partial t^2} \neq 0. \quad (4.10)$$

Also, the classic curvature scale-space image is extended from a binary image to a three-valued image to encode the concavity or convexity property of the curvature extrema. Local curvature extrema whose curvature value is greater than zero (i.e., concave curvature extrema) are marked as shaded points in the image, while local curvature extrema whose curvature value is less than 0 (i.e., convex curvature extrema) are marked as solid points. Such an extended three-valued scale-space image³³ for the smoothing process of Fig. 4.3 is depicted in Fig. 4.8a.

To construct the extended curvature scale-space image, the zero-crossings of the partial derivative of the curvature function with respect to t for each smoothing scale \mathbf{s} are determined and their positions are recorded on the image along with their concavity or convexity property. For example, the curvature extrema B, C and D at $\mathbf{s} = 48$ of Fig. 4.3 are recorded as corresponding points B, C, and D in Fig. 4.8a. The concavity property of C (as a shaded point) and the convexity property of B and D (as solid points) are also marked accordingly. The curvature function for $\mathbf{s} = 48$ and its partial derivative with respect to t are shown in Fig. 4.9.

³³ The extended curvature scale-space image is considered as a three-valued image because it requires 3 values to encode the concave curvature extrema, the convex curvature extrema and the background.






Figure 4.8. The extended and classic curvature scale-space images for Gaussian smoothing process shown in Fig. 4.3. (a) The extended curvature scale-space image. (b) The overlay of the classic and extended curvature scale-space image shown in Fig. 4.4 and Fig. 4.8a.



(a)



(b)

Figure 4.9. Curvature function and its first derivative for $\sigma = 48$ of Fig. 4.3. (a) The curvature function. The points B, D, N are convex curvature extrema, while C and M are concave curvature extrema. (b) The first partial derivative of the curvature function in (a) with respect to t . Curvature extrema in (a) become zero-crossing points.

4.3.5. Comparison of classic and extended curvature scale-space images

There are similarities and differences between the classic and the extended curvature scale-space processes. Both scale-space images contain many contour arcs which are formed by the loci of the respective investigated features. However, the important difference between the two processes is the selection of points from the curvature function of the smoothed borders. The classic scale-space process selects the curvature zero-crossing points while the extended scale-space process selects the curvature extrema points. Fig. 4.8b manifests the similarities and differences between the two scale-space images by overlaying Fig. 4.4 (the classic curvature scale-space image) and 4.8a (the extended curvature scale-space image). In this section, we present the parallel properties and the differences for these two images.

Property 1a: *In classic curvature scale-space images, the apex of a contour arc is the point (τ, ξ) such that $K(\tau, \xi)=0$ and $\partial K(\tau, \xi)/\partial t=0$.*³⁴

For any σ in the interval of $[0, \xi)$ in the classic curvature scale-space image, let the points t_1 and t_2 be the curvature zero-crossings at the two sides of the contour arc. Since $K(t_1, \sigma) = K(t_2, \sigma) = 0$ by definition (see Equation 4.9) and K is a continuous function, according to Rolle's Theorem, there exists a point t_3 such that $t_1 < t_3 < t_2$ and $\partial K(t_3, \sigma)/\partial t = 0$ in K - t space. At the smoothing level ξ , the points t_1 , t_2 and t_3 merge together to the point τ . Because K is continuous, $K(\tau, \xi)=0$ and $\partial K(\tau, \xi)/\partial t=0$.

For example, the points A2 and A3 in Fig. 4.3 and 4.4 are the curvature zero-crossing points for the smoothed border at $\sigma = 40$. The corresponding smoothed border is shown with its curvature function K_{40} in Fig. 4.10a and 4.10b, respectively. The point A_{Cl} is a concave curvature extremum between A2 and A3. At the contour apex (see Fig. 4.10c and 4.10d), the

³⁴ Note that the apex point (τ, ξ) of a contour arc is not selected in the classic curvature scale-space process due to the definition of the process as expressed in Equation 4.9. However, the property of the point can be derived.

loci of A2, A3 and A_{C1} merge together, and, therefore, the curvature of the merged point and its partial derivative with respect to t are equal to zero.

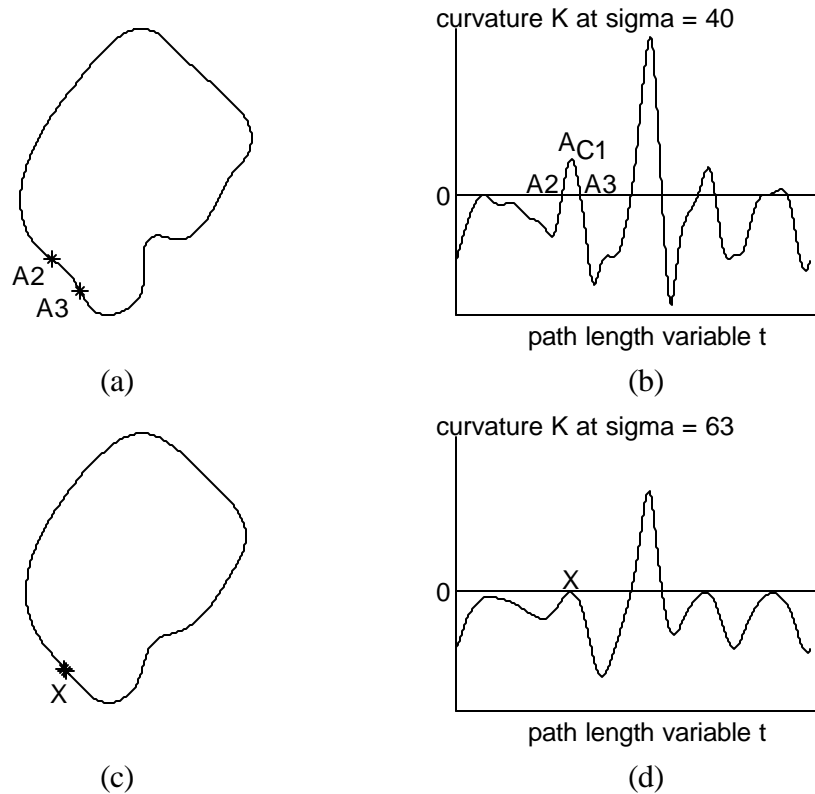


Figure 4.10. Smoothed borders of Fig. 4.3 at $\sigma = 40$ and $\sigma = 63$. (a) The smoothed border at $\sigma = 40$, the points A2 and A3 are curvature zero-crossing points. (b) The curvature function of (a) in $K-t$ space. The point A_{C1} is the concave curvature extreme which is between A2 and A3. (c) For the smoothed border at $\sigma = 63$, the point X marks the location where the loci of A2, A3 and A_{C1} merge together. (d) The curvature function of (c) in $K-t$ space.

Property 1b: *In extended curvature scale-space images, the apex of a contour arc is the point (τ, ξ) such that $\partial K(\tau, \xi)/\partial t=0$ and $\partial^2 K(\tau, \xi)/\partial t^2=0$.³⁵*

³⁵ Note that the apex point (τ, ξ) of a contour arc is not selected in the extended curvature scale-space process due to the definition of the process as expressed in Equation 4.10. However, the property of the point can be derived.

For any σ in the interval of $[0, \xi)$ in the extended curvature scale-space image, let the points t_1 and t_2 be the curvature extrema at the two sides of the contour arc. Since $\partial K(t_1, \sigma)/\partial t = \partial K(t_2, \sigma)/\partial t = 0$ by definition (see Equation 4.10) and $\partial K/\partial t$ is a continuous function, according to Rolle's Theorem, there exists a point t_3 such that $t_1 < t_3 < t_2$ and $\partial^2 K(t_3, \sigma)/\partial t^2 = 0$ in $\partial K/\partial t$ - t space. At the smoothing level ξ , the points t_1, t_2 and t_3 merge together to the point τ . Because $\partial K/\partial t$ is continuous, $\partial K(\tau, \xi)/\partial t = 0$ and $\partial^2 K(\tau, \xi)/\partial t^2 = 0$.

For example, the points D and M are curvature extrema in Fig. 4.3 and 4.8a. The corresponding smoothed border curve is shown along with its partial derivative of the curvature function with respect to t in Fig. 4.11a and 4.11b, respectively. There exists a local extremum A_{EI} such that $D < A_{EI} < M$. At the contour apex (see Fig. 4.11c and 4.11d), the loci of D, M and A_{EI} merge together, and, therefore, the first and second partial derivatives of K with respect to t are equal to zero at the merged point.

Property 2a: *In classic curvature scale-space images, excluding the apex point, one side of a contour arc has the property $\partial K/\partial t > 0$ and the other side of the contour arc has the property $\partial K/\partial t < 0$.*

Assume the contour apex is the point (τ, ξ) in the classic curvature scale-space image. For any σ in the interval of $[0, \xi)$ of the smoothing axis, let the points t_1 and t_2 be the curvature zero-crossings at the two sides of the contour arc. By definition, $K(t_1, \sigma) = 0$, $\partial K(t_1, \sigma)/\partial t \neq 0$ and $K(t_2, \sigma) = 0$, $\partial K(t_2, \sigma)/\partial t \neq 0$. (See Equation 4.9.) Let's first consider the point t_1 . Without loss of generality, we assume $K(t_1, \sigma) = 0$ and $\partial K(t_1, \sigma)/\partial t > 0$. In other words, for the smoothing level σ , K crosses the zero from below at t_1 in the K - t space. Because K is continuous, for K to cross the next zero at t_2 , K must cross the zero from above, i.e., $\partial K(t_2, \sigma)/\partial t < 0$. Otherwise, there exists a curvature zero-crossing in between t_1 and t_2 , which contradicts the classic curvature scale-space process. Therefore, $\partial K(t_1, \sigma)/\partial t$ and $\partial K(t_2, \sigma)/\partial t$ must have different sign.

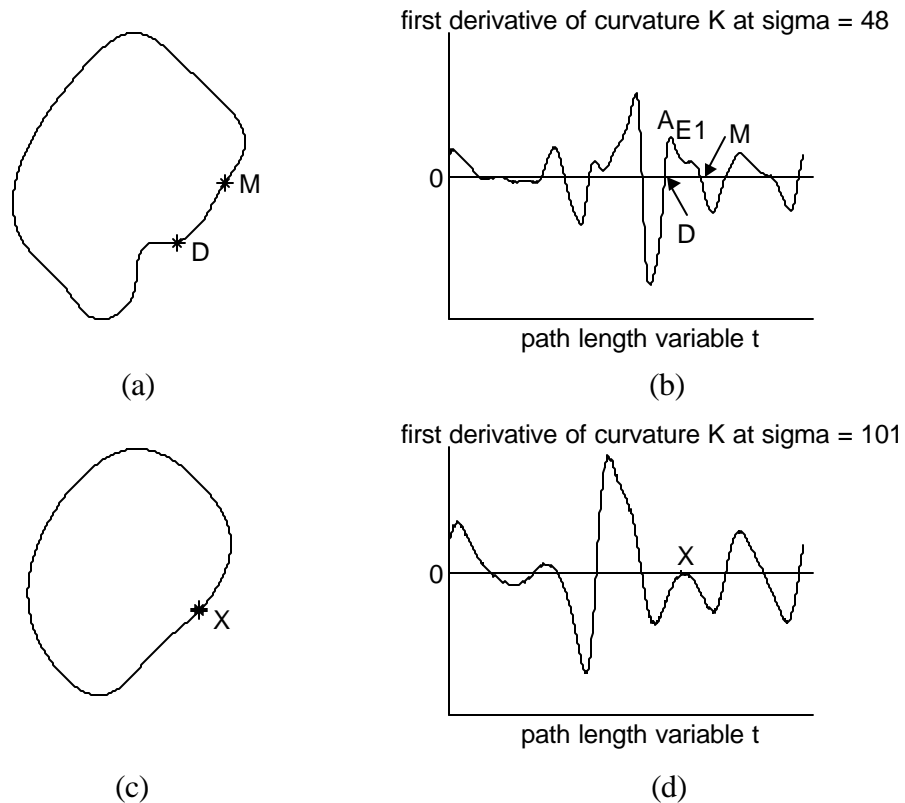


Figure 4.11. Smoothed borders of Fig. 4.3 at $\sigma = 48$ and $\sigma = 101$. (a) The smoothed border at $\sigma = 48$. The points D and M are curvature extrema. (b) The partial derivative of the curvature function K of (a) with respect to t in $\partial K/\partial t-t$ space. The point A_{EI} is a local maximum which is between D and M. (c) The smoothed border at $\sigma = 101$, the point X marks the location where the loci of D, M and A_{EI} merge together. (d) The partial derivative of the curvature function K of (c) with respect to t in $\partial K/\partial t-t$ space.

To complete our argument for the property, we have to show that if $\partial K(t_1, \sigma)/\partial t > 0$, all curvature zero-crossings in the same side of the contour arc must have the property $\partial K/\partial t > 0$. Since $\partial K(t_1, \sigma)/\partial t > 0$ and $\partial K/\partial t = 0$ only at the contour apex (τ, ζ) , moving along the contour arc from (t_1, σ) to (τ, ζ) in the $\partial K/\partial t$ surface cannot go to negative because $\partial K/\partial t$ is continuous. Therefore, the curvature zero-crossing along the same side as t_1 have the property $\partial K/\partial t > 0$.

Property 2b: *In extended curvature scale-space images, excluding the apex point, one side of a contour arc has the property $\partial^2 K/\partial t^2 > 0$ and the other side of the contour arc has the property $\partial^2 K/\partial t^2 < 0$.*

The argument is parallel to property 2a if we can show $\partial^2 K/\partial t^2$ is a continuous function. Since the border L_0 is C^2 , the smoothed border $L(t, \sigma)$ and curvature K are C^3 and $\partial^2 K/\partial t^2$ is C^1 . Therefore, $\partial^2 K/\partial t^2$ is a continuous function.

Assume the contour apex is the point (τ, ξ) in the extended curvature scale-space image. For any σ in the interval of $[0, \xi)$ of the smoothing axis, let the points t_1 and t_2 be the curvature extrema at the two sides of the contour arc. By definition, $\partial K(t_1, \sigma)/\partial t = 0$, $\partial^2 K(t_1, \sigma)/\partial t^2 \neq 0$ and $\partial K(t_2, \sigma)/\partial t = 0$, $\partial^2 K(t_2, \sigma)/\partial t^2 \neq 0$. (See Equation 4.10.) Let's first consider the point t_1 . Without loss of generality, we assume $\partial K(t_1, \sigma)/\partial t = 0$ and $\partial^2 K(t_1, \sigma)/\partial t^2 > 0$. In other words, for the smoothing level σ , ∂K crosses the zero from below at t_1 in the $\partial K/\partial t$ - t space. Because $\partial K/\partial t$ is continuous, for $\partial K/\partial t$ to cross the next zero at t_2 , $\partial K/\partial t$ must cross the zero from above, i.e., $\partial^2 K(t_2, \sigma)/\partial t^2 < 0$. Otherwise, there exists a curvature extrema in between t_1 and t_2 , which contradicts the extended curvature scale-space process. Therefore, $\partial^2 K(t_1, \sigma)/\partial t^2$ and $\partial^2 K(t_2, \sigma)/\partial t^2$ must have different sign.

To complete our argument for the property, we have to show that if $\partial^2 K(t_1, \sigma)/\partial t^2 > 0$, all curvature extrema in the same side of the contour arc must have the property $\partial^2 K/\partial t^2 > 0$. Since $\partial^2 K(t_1, \sigma)/\partial t^2 > 0$ and $\partial^2 K/\partial t^2 = 0$ only at the contour apex (τ, ξ) , moving along the contour arc from (t_1, σ) to (τ, ξ) in the $\partial^2 K/\partial t^2$ surface cannot go to negative because $\partial^2 K/\partial t^2$ is continuous. Therefore, the curvature extrema along the same side as t_1 have the property $\partial^2 K/\partial t^2 > 0$.

Property 3 *In the contours of an extended curvature scale-space image, the points where the concave extrema and convex extrema meet are the zero curvature points.*

The curvature of a convex curvature extremum is less than 0 and the curvature of a concave curvature extremum is greater than 0; hence, the meeting point has the property of zero curvature. The Points A_C and G in Fig. 4.8a and 4.8b are the examples of such points.

Even though the extended curvature scale-space process captures the locations of curvature extrema, some zero curvature points (the contour arc apices of the classic curvature scale-space image) can be identified easily with our 3-valued scale-space image. However, there is no corresponding property for the classic curvature scale-space image.

Property 4a: *In classic curvature scale-space images, all curvature zero-crossings disappear at σ_{term} .*

When a Gaussian smoothing process terminates at σ_{term} , the object is transformed into an oval shaped border with convex curvature for the entire border (i.e. $K(t, \sigma_{term}) < 0$ for all t); therefore, all curvature zero-crossings disappear.

Property 4b: *In extended curvature scale-space images, all curvature extrema may disappear (a special case of a circle) or at least 4 curvature extrema remain at σ_{term} .*

When a Gaussian smoothing process terminates at σ_{term} , the object is transformed into an oval shaped border with convex curvature for the entire border (i.e. $K(t, \sigma_{term}) < 0$ for all t). In a special case, $K(t, \sigma_{term})$ is a negative constant (i.e. a circle) and there will be no curvature extrema. Otherwise, curvature extrema must exist. Since an ellipse has 4 curvature extrema, there must be at least 4 curvature extrema remain at σ_{term} for the oval shaped border.³⁶

The most important difference between the two scale-space images is the functionality of the image. Classic curvature scale-space images are designed to analyse point features, while extended curvature scale-space images analyse indentation and protrusion curve segments,

³⁶ If the smoothing is carried out after σ_{term} , these convex curvature extrema may remain until the object border is turned into a point.

which can be organized into hierarchical structures. Details of the segment analysis are discussed in the following sections.

4.3.6. Identifying the span of indentation/protrusion segments

Our extended curvature scale-space image captures all indentation/protrusion segments as defined in Section 4.2 for the entire smoothing process. To reveal these segments at each smoothing level, we scan for three consecutive curvature extrema with alternating signs (shaded or solid points) sequentially from $\mathbf{s} = 0$, the original non-smoothed curve, to \mathbf{s}_{term} .

As an indentation/protrusion segment evolves through the smoothing process, the segment may span several scales. Unfortunately, Gaussian smoothing distorts the contour length and, hence, the location of indentation/protrusion segments shifts from scale to scale. Matching up the segments between scales becomes a difficult task. The extended scale-space image facilitates the matching by analysing the loci formed by the curvature extrema points. For example, tracking the convex extrema lines B and D of the indentation segment [B, C, D] at $\mathbf{s} = 48$ of Fig. 4.8a down toward the zero-scale reveals the cover, the *true* position of the segment specified by the non-smoothed curve at $\mathbf{s} = 0$. More precisely, for any segment $U = [t_1, t_2, t_3]$ at smoothing level \mathbf{s} , the cover Γ_U specifies the segment's corresponding position at the zero-scale and is defined as:

$$\Gamma_U = [u_1, u_2], \quad (4.11)$$

where u_1 and u_2 are the path length variables for the beginning and ending positions of the segment. The positions u_1 and u_2 are obtained by coarse-to-fine tracing the loci of the terminal curvature extrema t_1 and t_3 . Γ_U is considered the *true* position of the segment U , and hence, matching segments found in different scales can be done by comparing their corresponding covers. For any two segments U and V at the levels \mathbf{s}_U and \mathbf{s}_V , respectively, where $\mathbf{s}_U > \mathbf{s}_V$, segments U and V are considered the same segment if $\Gamma_U = \Gamma_V$.

Identifying the span of a segment in the smoothing scale axis reveals the evolution of the segment. There are two important properties for a segment evolution: the formation level and

the smooth-out level. The formation level indicates the first appearance of an irregularity segment, while the smooth-out level indicates the disappearance of the segment. More precisely, let the maximum sigma level spanning for a segment U with the cover $\Gamma_U = [u_1, u_2]$ be $i, i+1, \dots, j$. In other words, for the consecutive sigma levels ranging from $i, i+1, \dots, j$, there exists an irregularity segment whose cover is $[u_1, u_2]$ in each level and there is no irregularity segment with cover $[u_1, u_2]$ in the sigma level $i-1$ and $j+1$. Then the formation level of the irregularity segment U is defined as i and the smooth-out level is defined as $j+1$. For example, the indentation segment [B, C, D] emerges at $\mathbf{s} = 44$ in Fig. 4.8a when all nested smaller irregularities have been smoothed out, and it ends at $\mathbf{s} = 102$ when the convex curvature line D closes off at the top with the line E.

4.3.7. Hierarchical structures for indentation/protrusion segments

The extended scale-space image not only illustrates the evolution of the indentation/protrusion segments; it also helps organize segments into a hierarchical structure. Because of the causality property of Gaussian smoothing [94, 104], segments are smoothed out in a 'proper' order: small ones disappear before larger ones. Now when some smaller segments are smoothed out, larger segments may emerge at the same locations. The larger segments are considered as the 'global' segments to the smaller 'local' ones. Hence, a hierarchical structure of indentation/protrusion segments is formed.

As illustrated by Fig. 4.2, adjacent indentation segments and protrusion segments may overlap at a smoothing level. To avoid the complexity of overlapping segments within a hierarchical structure, we divide indentation segments and protrusion segments into two separate hierarchies, which are constructed by examining the inclusion property of the segment covers. For a segment U with the cover $\Gamma_U = [u_1, u_2]$ and the smooth-out level \mathbf{s}_U and a segment V with the cover $\Gamma_V = [v_1, v_2]$ and the smooth-out level \mathbf{s}_V , segment U is a local segment for the global segment V if $\mathbf{s}_U < \mathbf{s}_V$ and Γ_U is included inside Γ_V . (See Fig. 4.12.) The inclusion of Γ_U inside Γ_V , denoted by $\Gamma_U \subseteq \Gamma_V$, is expressed as:

$$u_1 \geq v_1 \text{ and } u_2 \leq v_2. \quad (4.12)$$

The segments U and V form a hierarchical tree structure because a global segment can nest one or more local segments. For instance, the indentation segments $[B, C, D]$ and $[D, M, N]$ are local indentation segments for the global segment $[B1, C1, F]$ shown in Fig. 4.8a. (These segments are also shown in Fig. 4.3 at $s = 48$ and 102 , respectively.³⁷) By analysing the inclusion property of all indentation and protrusion segments, two hierarchical structures are obtained. One is for indentation segments and the other is for protrusion segments. Each hierarchical structure consists of a forest of tree structures, where the root of a tree structure is a global segment and its corresponding local segments are in the leaves of the tree.




Figure 4.12. Hierarchical relationship between segments. Segments are abstracted by their covers and are plotted at the smooth-out level, which is represented by the y-axis of the plot. Segment U with covers $[u_1, u_2]$ is included inside the segment V whose cover is $[v_1, v_2]$. Furthermore, segment R and S with covers $[r_1, r_2]$ and $[s_1, v_2]$, respectively, are also included inside segment V . Note: A segment can wrap around the beginning point of the parameterization process such as segment W , whose cover is $[w_1, w_2]$, where $w_1 > w_2$.

³⁷ Actually, the indentation segment $[D, M, N]$ is best illustrated at $s = 32$ as it is almost smoothed out at $s = 48$.

Note that all smoothed borders are periodic. Special care has to be taken when a segment wraps around the beginning point of the parameterization process such as the protrusion segment [H, I, J] shown in Fig. 4.8a and in Fig. 4.3, $\mathbf{s} = 32$. (See also segment W in Fig. 4.12.) The hierarchical relation between segments need to be checked carefully.

To properly analyse the inclusion property of a segment U with the cover $\Gamma_U = [u_1, u_2]$ and the smooth-out level \mathbf{s}_U , and a segment V with the cover $\Gamma_V = [v_1, v_2]$ and the smooth-out level \mathbf{s}_V , where $\mathbf{s}_U < \mathbf{s}_V$, we have to consider 4 sub-cases depending on whether segment U and/or segment V wrap around the beginning point of the parameterization process.

Sub-case 1: Both the segment U and V do not wrap around. Equation 4.12 can be applied directly for verifying the inclusion property of Γ_U inside Γ_V .

Sub-case 2: Both the segment U and V wrap around the beginning point of the parameterization process, i.e., $u_1 > u_2$ and $v_1 > v_2$. The segment U can be divided in the sub-segments $[0, u_2]$ and $[u_1, 1]$ and the segment V can be divided into $[0, v_2]$ and $[v_1, 1]$. The inclusion property of Γ_U inside Γ_V can be expressed as:

$$(u_2 \leq v_2) \quad \text{and} \quad (u_1 \geq v_1). \quad (4.13)$$

Sub-case 3: Segment U does not wrap around, but segment V does, i.e., $v_1 > v_2$. The segment V can be divided into $[0, v_2]$ and $[v_1, 1]$. The inclusion property of Γ_U inside Γ_V can be expressed as:

$$(u_1 \geq 0 \text{ and } u_2 \leq v_2) \quad \text{or} \quad (u_1 \geq v_1 \text{ and } u_2 \leq 1). \quad (4.14)$$

Sub-case 4: Segment U wraps around (i.e., $u_1 > u_2$), but segment V does not. Since segment U spans both sides of the beginning point of the parameterization process and segment V occupies only one side, segment U cannot be included inside segment V .

4.3.8. Flat irregularity segments

As local segments smooth into global segments and the border turns into an oval shape, the overall curvature of the curve is reduced. Global segments tend to be flatter than their

counterparts, local segments. Without a restriction on the formation of a segment, sometimes a flat segment, which is a slightly bent indentation/protrusion segment, can be formed. This kind of flat segment often occurs near the root position of a tree structure in a segment hierarchy. For example, when the indentation segment [B, C, D] in Fig. 4.3 is smoothed out at $s = 102$ and turns into a larger indentation segment [B1, C1, F], the new global segment [B1, C1, F] is a flat segment with a hardly noticeable indentation in the middle of the segment. (The indentation segments [B, C, D] and [B1, C1, F] are also labelled in the extended scale-space image in Fig 4.8a.) In order to remove the formation of such a 'flat' irregularity segment, the three curvature extrema are checked at the formation level for every new segment $U = [t_1, t_2, t_3]$. If the absolute magnitude of the middle curvature extremum, t_2 , or the maximum absolute magnitude of the first and the last curvature extrema, t_1 , and t_3 , are smaller than certain threshold and very close to zero, the newly formed segment is considered as an insignificant 'flat' indentation/protrusion segment. Mathematically, the criteria for a flat segment can be written as:

$$|K(t_2)| < \mathbf{e} \quad \text{or} \quad \max\{|K(t_1)|, |K(t_3)|\} < \mathbf{e}, \quad (4.15)$$

where \mathbf{e} is the threshold, which is set to 0.01. All flat segments are removed from further computation.

In summary, by analyzing the extended scale-space image, we can track the evolution of all irregularity segments for the entire smoothing process. These segments can span multiple smoothing levels and they can be organized separately in two hierarchical structures. Furthermore, the smooth-out level for each segment is an important piece of information that can be used in the computation of the border irregularity.

4.4. Calculating irregularity indices

Each detected indentation/protrusion segment has an irregularity measure, which is formed by observing the smoothing effect on the area of an indentation or a protrusion segment. When an indentation (or a protrusion) is smoothed-out, the indentation (or protrusion) is partially filled (or removed). For example, Fig. 4.13a shows a lesion border and the smoothed contour at the

smooth-out level for the largest indentation at the bottom of the figure. The shaded area indicates the filling done by the smoothing process. Likewise, Fig. 4.13b shows the same lesion border and the smoothed border for the most prominent protrusion at the bottom of the border. The shaded area represents the area removed by the smoothing process. The size of the filled (or removed) area, termed irregularity area, is used to determine the irregularity index.




Figure 4.13. Smoothing effect on indentation and protrusion. (a) A lesion border is shown by the solid line, while the smoothed curve corresponding to the smooth-out σ level for the largest indentation is shown by the dashed line. The shaded area denotes the irregularity area filled by the smoothing process. (b) The smoothed curve corresponding to the smooth-out σ level for the largest protrusion is shown by the dashed line. The irregularity area removed by the smoothing process is marked by the shaded area.

The index for an indentation/protrusion segment must be normalized so that it can be used for comparison among irregularities in different lesion borders. Normalization is achieved by dividing the irregularity area by the area of smooth-out contour, which is the area of the smoothed border at the smooth-out σ level. For example, the dashed line in Fig. 4.13a and 4.13b denote the smooth-out contour for the largest indentation and protrusion. The areas enclosed by the dashed lines are the normalization factors for the corresponding irregularity segments. Therefore, the Irregularity Index (II) of an indentation/protrusion segment U , denoted by I_U , is defined as:

$$I_U = \frac{\Delta_U}{R_U} \times 100\%, \quad (4.16)$$

where Δ_U is the irregularity area for segment U and R_U is the area enclosed by the contour of the smooth-out level. Because R_U is proportional to the original lesion size, I_U can be considered as the ratio of the irregularity area and the lesion size. As the ratio increases, the irregularity segment appears more prominent and, hence, the II increases.

The hierarchical structures described in Section 4.3.7 organize all indentation and protrusion segments of the entire smoothing process into two forests of tree structures. In particular, the set of root segments in the tree structures of both hierarchies represents all global irregular segments, which fully describe the complexity of the lesion border. Let's assume the root segments are U_1, U_2, \dots, U_n . Their corresponding II 's, denoted by I_1, I_2, \dots, I_n respectively, provide a rich set of descriptions of the border. From this set of indices, many important parameters about a lesion border can be inferred. In particular, two important descriptors, the Most Significant Irregularity Index (MSII) and the Overall Irregularity Index (OII), can be derived. The MSII of a lesion border L ranks all individual indices and indicates the largest indentation/protrusion segment of the border:

$$MSII = \max\{I_1, I_2, \dots, I_n\}. \quad (4.17)$$

The OII represents the entire lesion shape, and is calculated by summing up all individual indices. Thus:

$$OII = \sum_{j=1}^n I_j, \quad (4.18)$$

Fig. 4.14 depicts ten largest indentation/protrusion segments for the lesion border shown in Fig. 4.1. The segments are sorted by their corresponding II . The top left-hand subfigure depicts the largest irregularity segment, a protrusion, with the $MSII = 4.2$, while the next subfigure illustrates the largest indentation segment with an index of 2.4. Note that the indentation segment overlaps partially with the largest protrusion segment of the border. The

third largest irregularity segment is another protrusion with an index of 1.9. The OII for this lesion border is 15.1.




Figure 4.14. Indentation/protrusion segments of Fig. 4.1. Ten largest global indentation and protrusion segments for lesion border shown in Fig. 4.1 are plotted. The segments are sorted by their irregularity index, which is shown at the top of each subfigure.

4.5. Sensitivity analysis of irregularity indices

The new area-based method overcomes the non-linearity and the long-and-narrow indentation problem with the earlier SR method reported in Section 4.3.3. The MSII and the OII for the phantom shown in Fig. 4.7a (also shown in Fig. 4.15a) are 0.3 and 0.4, respectively.³⁸ They differ significantly from the MSII (7.4) and the OII (7.5) for the phantom in Fig. 4.7b (also shown in Fig. 4.15b). (The comparison of the SR, OII and MSII for these two phantoms are summarized in Table 4.1.) The small MSII value for Fig. 4.15a reflects the small area of the indentation and implies the stability of the new method. Small alterations (noise) are detected as

³⁸ The small texture irregularity, 0.1, is due to the discrete representation of a continuous curve.

local irregular segments with small irregularity areas, which do not significantly alter the final analysis results.

Table 4.1. Comparison of the SR, MSII and OII for two phantoms .

	Phantom shown in Fig. 4.15a	Phantom shown in Fig. 4.15b
SR	0.96	0.97
MSII	0.3	7.4
OII	0.4	7.5




Figure 4.15. The most significant indentations for two phantoms shown in Fig. 4.7. (a) and (b) The most significant indentation is highlighted and the associated MSII is shown at the top of each phantom.

4.6. The advantages of irregularity indices

As discussed in Section 2.1, many measures have been used to estimate the border irregularity of a melanocytic lesion. However, each of these methods has some drawbacks. The CI, form factor, FD, and SFD are insensitive to structure indentations and protrusions. The convex hull method fails to account for protrusion, one of the important clinical features for the lesion. Many methods are sensitive to noise along the border. For example, the CI amplifies the noise by a square term. The radial distance distribution method and the circularity factor may be

unstable due to their dependence of the centroid location, which is very sensitive to noise. Furthermore, all these methods estimate the border shape using a single value. Lesions with completely different shapes may associate with the same value.

Our irregularity index method addresses the above problems by detecting all global and local irregularity segments along the border. For each irregular segment, an area-based index is computed. Then an overall irregularity index and the most significant irregularity index are derived from the set of indices. This approach has many advantages:

1. Our method detects both indentation and protrusion segments, which are important clinical features for melanocytic lesions.
2. Our method is sensitive to structure irregularities because the multi-scale method actually locate all local and global segments in a rugged border and organize the segments into hierarchical structures.
3. Our method is stable. The area-based approach implies the method is stable for small changes (noise) along the border.
4. Our method is linear. Defined as the sum of all the irregularity indices of the global irregularity segments, the OII is proportional to the total irregularity areas of the lesion.

Our method returns a set of irregularity segments with corresponding indices, which fully describe the complexity of the lesion shape. This information set can be used to derive many parameters for the lesion. In particular, the overall irregularity index and the most significant irregularity index are two important derived indices.

Chapter 5

A User Study

To verify the effectiveness of the two new measures (the MSII and OII), a user study has been carried out to compare the new measures as well as other shape descriptors, namely the CI, the FD and the SFD, with 20 dermatologists' clinical evaluations.³⁹ The CI was selected because it is the most common shape descriptor based on single-scale computation, while the FD and the SFD represent commonly used multi-scale methods. Forty lesion borders selected from the B.C. Cancer Agency pigmented lesion image database were used in the study. These tested measurements were analysed statistically.

5.1. Method

5.1.1 Gold standard

One of the considerations for designing the user study was the gold standard selection for the validation process. Should it be the clinical evaluation or histology status? Because border irregularity is a clinical diagnosis feature defined by dermatologists, making dermatologist's clinical evaluation as the gold standard was the most appropriate choice. In other words, when dermatologists had concerns on an irregular border, the new measures should reflect the similar concerns. A computer program encoding such knowledge could be used by other non-dermatologists such as health practitioners or the general public. Also it could be used as an objective alternative for dermatologists.

5.1.2. Assembling the data

Forty pigmented lesions were selected from the skin image database. In order to have a good representation of lesions in the entire range of the OII (one of the new measures to be verified)

³⁹ A user study with 3 dermatologists has been reported earlier [89].

and to void possible biases, the following selection scheme was employed. Many melanocytic lesion images were first processed by DullRazor and the automatic segmentation programs described in Section 3.2. 108 of these melanocytic lesions with good segmentation results were retained in the selection process. The OII of these lesions were computed, and the sorted OII values are plotted in Fig. 5.1. All OII values increased smoothly, except the 4 outliers (marked as O) in the right-hand side of Fig. 5.1. These 4 outliers were included in the final selection for the user study data set. For the rest of the 104 lesions, the 10 lesion borders with the lowest OII values (denoted as ∇), the 10 lesion borders with the highest OII values (denoted as Δ) and the middle 16 lesions (denoted as $+$) were also included. In total, we chose 40 lesion borders, spanning the entire range of the OII values. These lesion borders are depicted in Fig. 5.2.




Figure 5.1. The overall irregularity index (OII) of 108 lesions. The 40 lesions marked as ∇ , $+$, Δ and O were selected for the user study.




Figure 5.2. Forty lesion borders selected for the user study.

5.1.3. The compactness index implementation

The CI was implemented according to Equation 2.1. The program was tested by real-valued circles of various sizes and the resultant values were 1.00, identical to the theoretical value. However, when the circle was converted to a discrete circle with a radius 150-pixel by rounding the x and y coordinates to the nearest integer, the derived CI value became 1.11 (11% error) because the discrete contour amplified the error.

5.1.4. The fractal dimension implementation

The FD of a lesion border L , denoted by D in Equation 2.3, was estimated using the box counting method [28, 108]. The border L was represented by a binary mask where the border was specified by 1 and the background was specified by 0. To ensure there were no partial blocks, the background was extended to the nearest integer with power of 2. Various sizes of $r \times r$ boxes, where r was power of 2, were placed over the lesion and the number of boxes containing the border, $N(r)$, was counted. Applying r and $N(r)$ to Equation 2.3, the value of D could be estimated. Rearranging the terms and expanding Equation 2.3 by taking the log, the equation could be expressed as:

$$\log \frac{1}{N(r)} = D \times \log(r) - \log(I), \quad (5.1)$$

where I was a constant. Thus, D was the slope of the linear equation, which could be computed by least square fitting of $\log(1/N(r))$ vs. $\log(r)$. Fig. 5.3 gives a sample of the log-log plot using the lesion border of Fig. 4.1.




Figure 5.3. An example of the log-log plot for fractal dimensions. This plot was constructed using the lesion border shown in Fig. 4.1. The FD is the slope of the solid line, estimated by all the data points. The SFD is the slope of the dashed line (- -), estimated by the fifth to the ninth data points. The FD for this lesion border is 1.12, and the SFD is 1.21.

When a discrete circle with a radius of 150 pixels (the x and y coordinates were rounded to the nearest integers) was used to test the FD implemented, the resultant FD was 0.99. (The error was 1% as the theoretical value of a circle is 1.) Further tests were done by using the well-known Koch snowflake (see Fig. 5.4a) and Koch square flake (see Fig. 5.5a) [117]. The Koch snowflake was constructed using an equilateral triangle. Each side of the triangle was transformed recursively by dividing the line segment into 3 equal parts and replacing the middle part with an equilateral triangle without a base. Fig. 5.4b illustrates the construction of the Koch snowflake. The result of the transformation has a theoretical FD value of $\log 4 / \log 3 \cong 1.26$. The estimated FD from the implementation was 1.29 (2.38% error). The Koch square

flake was constructed by 4 equal line segments placed as a square. Each segment was transformed recursively by dividing the segment into 4 equal parts and replacing the middle two parts by two squares with no bases at the reverse orientations as illustrated in Fig. 5.5b. The theoretical FD for Koch square flake is 1.5 and the computed FD was 1.50 (0% error).

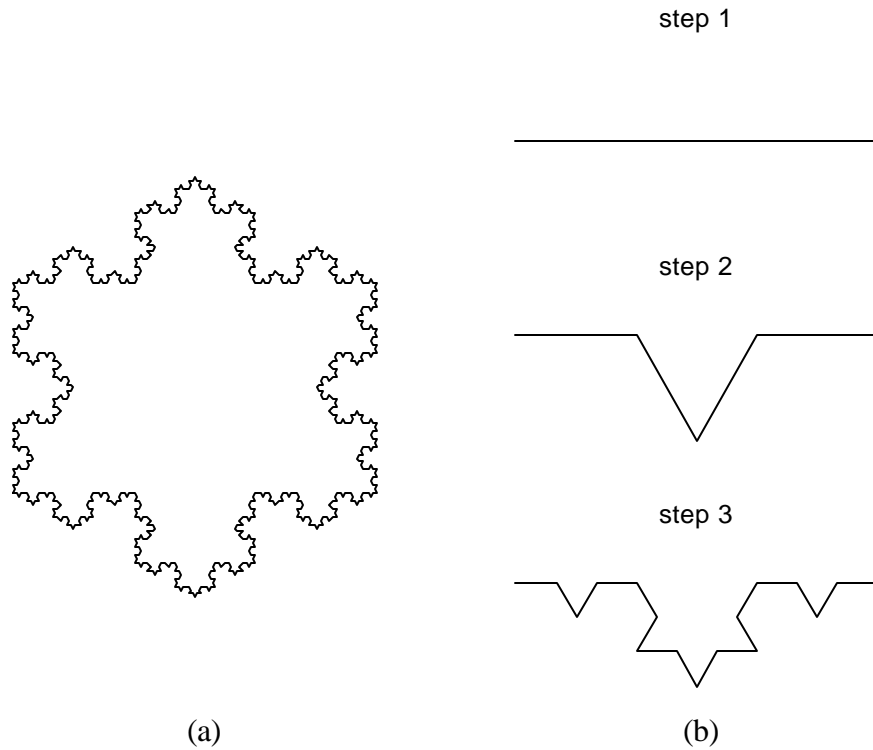


Figure 5.4. Koch snowflake and its construction procedure. (a) Koch snowflake. (b) The construction procedure: at each step, the middle section of a line segment is replaced by an equilateral triangle with no base.




Figure 5.5. Koch square flake and its construction procedure. (a) Koch square flake. (b) The construction procedure: at each step, the middle two sections of a line segment are replaced by two square with no bases at the reverse directions.

As mentioned in Section 2.1, Claridge *et al.* [31] recognized that there were two types of fractal dimensions, structure and texture fractal dimension, associated with a lesion border. Furthermore, the SFD had a higher correlation with melanomas. When the log-log plots for the fractal dimensions were examined (see Fig. 5.3 as an example for the log-log plot), evidence to support the idea of two fractals was found. The breakpoint for these two fractals was estimated to be located between the fourth and fifth data points in the log-log plot. Thus, the data points were divided at the breakpoint and the upper half of data, from the fifth data point, was used to compute the SFD of the lesion border. The SFD's of a discrete circle with a radius of 150 (the x and y coordinates were rounded to the nearest integers), Koch snowflake and Koch square

flake were 1.05 (5% error), 1.29 (2.38% error) and 1.50 (0% error), respectively. The errors for both the FD and SFD were considered small enough to be acceptable.

5.1.5. The irregularity index implementation

The methodological development described in Chapter 4 has been implemented in Matlab, running on a Pentium 233MHz computer. The smoothing s level was incremented by unit steps, which meant the computation was extensive. The execution time ranged from several minutes to an hour, depending on the border length.⁴⁰ For each tested lesion border, the MSII and OII were computed. When a discrete circle of radius 150 pixels (the x and y coordinates were rounded to the nearest integers) was input to the program, the OII and MSII were 0.06 (absolute error = 0.06) and 0.00 (absolute error = 0.00), respectively. The theoretical values should be zero for both indices, but the computational error was insignificant.⁴¹

5.1.6. Clinical evaluation

Twenty experienced dermatologists were asked to visually evaluate the 40 tested lesion borders shown in Fig. 5.2 using a scoring scale of 1 to 4. A user study package was prepared by plotting the lesion borders into 4 pages, 10 borders per page. Because the size of a lesion is itself a marker of risks, all lesion sizes were standardized to an equal area so that the evaluation was based solely on the border features. To avoid all lesion borders with high (or low) OII values clustering into one page, the following placement scheme was employed. The four lesion borders associated with the highest OII values (the four outliers marked as O in Fig. 5.1) were placed separately, one in a page. The rest of the 36 lesion borders were divided into 4 groups according to their OII values and each group was randomized and distributed evenly in the four pages. Furthermore, within a page, the placement of a lesion border was randomized. The

⁴⁰ The computation time can be improved substantially by using a faster computer language such as C, skipping some smoothing s levels especially at the coarse scales, and, of course, running the program on a faster computer.

⁴¹ The percentage of error cannot be computed as the based value (theoretical value) is 0.

dermatologists were informed only that there was at least one border in the scale category 1 and at least one border in the scale category 4 in each page. However, the number of lesions in each of the 4 scale categories in a page and in the entire study were not mentioned. The dermatologists could freely assign a lesion border into any scale category and any number of borders into a category. A user consent form and an information sheet were also included in the package. (See Appendix A for the entire user study package.) The information sheet described the study and provided evaluation instructions. The dermatologists were asked to read over the instructions and complete the evaluation. Other than the study package, there was no further communication with the dermatologists. The study design and material were approved by the University Ethics Review Committee of Simon Fraser University (SFU).

The evaluation was done on a scale of 1 to 4. The smallest scale 1 implied the most benign looking border contour and the highest scale 4 implied the most severe case with the highest probability of being a melanoma. This was a double-blinded test because all program development and calculations were done before the clinical evaluation and the dermatologists did not know the results of the tested measurements prior to the evaluation. The computed measurements were then compared statistically against the average of the 20 dermatologists' evaluation, which was considered as the gold standard.

5.1.7. Statistical methods

The user study results were tested by statistical methods using SPSS [109]. First the Kendall W , the coefficient of concordance, was computed for the dermatologists' evaluation to ensure a reasonable agreement among the dermatologists so that the gold standard could be formed. The null hypothesis of the test, often denoted as H_0 , is that there was no agreement among the ratings of k judges (20 dermatologists) on n objects (40 lesion borders). In such a situation, after the ratings for each judge is ranked from 1 to n , the sum of each object rankings among all the judges would be approximately equal to the sum of average ranking, $k(n+1)/2$. The difference between the sum of an object ranking among all judges and the sum of average ranking exhibits the amount of deviation from H_0 . Kendall W is defined by the sum of the

square of such a difference for all objects with the appropriate normalization. The value of W ranges from 0 to 1, where 0 implies there is no agreement among the judges and 1 implies a perfect agreement. The significance of W is indicated by a probability p , where a large p ($p > 0.05$) implies that H_0 cannot be rejected. However, a small p , where $p \leq 0.05$, rejects the null hypothesis and accepts the alternative hypothesis, often called H_1 , that there is an agreement among the judges' rating [32, 109, 135]. The gold standard should be formed by averaging the clinical evaluation for each tested lesion only with a good agreement among the dermatologists (with a large W and a small p value).

Then the Spearman's rank correlation coefficient was determined for each pair of measurements, the average clinical evaluation, CI, FD, SFD, OII and MSII, using SPSS [109]. The Spearman coefficient is the well-known Pearson correlation based on the rank of the measurements. Its value ranges from -1 to +1, where -1 (+1) implies a perfect negative (positive) correlation between the rank of two measurements, and 0 implies there is no linear correlation between the rank of two measurements.

Finally, the relationship between the average clinical evaluation and the CI, FD, SFD, OII and MSII was tested out using multiple linear regression analysis under SPSS [109]. The average clinical evaluation was set to be the dependent variable Y of the linear regression model and all tested measurements (CI, FD, SFD, OII and MSII) were the possible independent variables, which were selected in the stepwise fashion. The tested measurement that has the best linear relationship with Y was entered into the model as an independent variable. Then the linear relationship between Y and all dependent variables in the model were examined to ensure their linear relationships were not altered due to the new entry. Any independent variable that did not demonstrate a linear relationship any longer was removed. The procedure was repeated until there were no more entries or removals. The final independent variables in the model were the best predictors of Y , the average clinical evaluation.

5.2. Results

The user study results are presented in this section and the details of the results are discussed in the next section. The clinical evaluation results of 20 dermatologists are reported in columns "L1" to "L20" in Table 5.1. The coefficient of concordance, Kendall W, for all 20 clinical evaluations was then determined as $W = 0.77$ ($p = 0.000$). The average of the clinical evaluations is listed in column "AvgCL" of Table 5.2. The tested measurements (CI, FD, SFD, OII, and MSII) were computed and presented in Table 5.2. The melanomas (based on pathological assessments) are indicated in column "Path" of Table 5.2. The most significant indentation/protrusion segment for each tested border is plotted in Fig. 5.6. The Spearman coefficients, shown in Table 5.3, were calculated for each pair of the average clinical evaluation and the tested measurement. Among all tested measurements, the OII achieved the highest correlation coefficient (0.88) with the average clinical evaluation. To further expose the relationship between the average clinical evaluation and the tested measurements, the tested measures are plotted against the average clinical evaluation separately in Fig. 5.7. Finally, the multiple linear regression analysis result is presented in Table 5.4. After the stepwise regression analysis, only one independent variable OII was left in the model.

5.3. Discussion

5.3.1. Clinical evaluation and gold standard

Clinical evaluation of border irregularity is a difficult task for dermatologists as they are trained to diagnose melanocytic lesions using not only border shapes, but also other features including lesion colour and size. Even experienced dermatologists have difficulty when interpreting the lesion border as the sole feature. Claridge *et al.* [32] reported a low agreement (coefficient of concordance $W = 0.47$) among 20 clinicians, when they were asked to sort 20 lesion borders contours in the order of increasing border irregularity. Such a low agreement cast doubt on a possible valid gold standard. Therefore, forming the gold standard becomes an issue.

Table 5.1. Clinical evaluation results of the lesion borders shown in Fig. 5.2

Border	L1	L2	L3	L4	L5	L6	L7	L8	L9	L10	L11	L12	L13	L14	L15	L16	L17	L18	L19	L20
1	2	1	2	1	2	2	2	1	1	1	2	1	2	2	1	2	2	1	1	2
2	3	2	2	1	2	2	2	1	1	1	2	2	3	2	2	3	2	1	2	2
3	2	3	4	2	3	3	2	2	3	3	3	3	4	3	3	3	3	2	3	3
4	4	4	3	4	4	4	4	3	4	4	4	4	4	4	4	4	4	4	4	4
5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2	1	1	1	1
6	1	1	2	1	2	2	2	2	1	1	2	1	2	2	2	1	2	1	1	2
7	3	3	4	4	4	3	3	3	3	2	3	3	4	2	3	3	3	3	4	3
8	2	1	1	1	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1
9	3	2	3	3	3	2	2	2	2	2	2	2	3	1	2	3	3	1	2	2
10	3	4	4	4	4	4	3	4	3	4	3	3	4	3	4	3	3	4	4	4
11	2	2	3	1	2	2	2	1	1	2	1	1	3	1	2	3	2	1	2	1
12	1	1	2	1	2	2	2	1	1	1	2	1	1	1	2	2	2	1	1	1
13	2	3	2	1	2	2	1	2	2	2	2	2	2	1	2	1	1	1	1	2
14	3	3	2	2	3	3	2	3	2	3	2	4	3	2	3	3	3	2	2	2
15	1	1	3	1	1	1	2	1	1	1	1	1	2	1	1	1	1	1	2	1
16	1	2	2	2	2	2	2	1	1	1	1	2	1	1	1	2	2	1	1	1
17	2	1	1	1	1	1	2	1	1	1	1	1	1	1	1	2	2	1	1	1
18	4	4	4	3	4	3	3	3	3	4	3	4	4	4	4	4	3	4	4	3
19	4	4	4	4	4	3	4	4	4	4	4	4	4	2	4	4	4	4	4	4
20	2	3	4	4	4	4	3	3	2	3	3	3	4	3	4	3	3	3	4	3
21	3	4	4	3	3	3	3	3	3	4	4	3	4	2	3	4	3	3	4	4
22	1	2	2	1	3	2	2	2	2	2	2	2	2	2	2	1	2	1	2	2
23	2	3	3	2	2	2	2	3	2	3	3	3	4	2	2	2	4	2	3	3
24	1	2	3	2	3	2	3	2	1	2	2	3	2	3	2	1	2	2	3	2
25	2	1	2	1	2	3	2	2	1	2	2	3	2	3	2	1	2	2	2	2
26	3	2	2	1	2	2	2	2	1	2	2	2	1	1	2	2	2	1	1	1
27	2	1	2	1	2	1	2	1	1	1	1	1	1	1	1	1	1	1	1	1
28	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4
29	1	1	1	1	1	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1
30	3	2	2	1	2	2	2	2	1	2	2	1	2	1	1	2	2	1	2	2
31	1	1	3	1	2	1	2	1	1	2	2	1	1	1	3	1	2	1	1	1
32	1	3	4	2	2	1	3	2	2	3	2	2	3	2	2	2	3	2	2	3
33	2	3	3	2	3	1	3	2	2	2	2	2	3	2	3	2	3	2	2	3
34	3	3	4	2	3	2	3	2	2	3	3	3	4	3	3	3	3	3	3	2
35	2	2	2	1	2	1	2	1	2	2	1	2	2	1	1	2	2	1	2	1
36	2	2	2	1	1	2	2	1	1	2	1	2	1	1	2	2	2	2	2	1
37	1	1	2	1	1	1	1	1	1	1	1	1	1	1	1	2	1	1	1	1
38	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4
39	2	3	3	2	2	3	2	3	2	2	2	3	4	2	3	3	3	3	2	2
40	1	1	1	1	1	2	2	1	1	2	1	1	1	1	1	2	2	1	1	1

Table 5.2. Average clinical evaluation and tested measurements for lesion borders shown in Fig. 5.2. The 4 melanomas, based on pathological assessments, are indicated in column "Path".

Border	AvgCL	CI	FD	SFD	OII	MSII	Path
1	1.6	3.44	1.18	1.35	5.7	1.1	-
2	1.9	3.75	1.12	1.22	5.5	0.6	-
3	2.9	2.30	1.08	1.20	8.7	1.4	-
4	3.9	3.44	1.12	1.21	17.1	3.2	melanoma
5	1.1	1.69	1.10	1.24	2.8	0.4	-
6	1.6	1.90	1.09	1.26	3.5	0.4	melanoma
7	3.2	3.72	1.15	1.24	5.4	1.1	-
8	1.2	3.07	1.10	1.12	3.4	0.2	-
9	2.3	2.13	1.08	1.22	5.7	1.0	-
10	3.6	3.90	1.13	1.25	8.2	2.2	-
11	1.8	2.32	1.08	1.12	5.1	0.7	-
12	1.4	2.05	1.08	1.18	5.2	1.3	-
13	1.7	3.19	1.13	1.22	5.0	0.7	-
14	2.6	4.95	1.18	1.28	9.4	1.4	-
15	1.3	1.49	1.11	1.40	3.4	0.8	-
16	1.5	2.26	1.11	1.26	5.7	1.7	-
17	1.2	3.02	1.18	1.35	3.3	0.3	-
18	3.6	4.50	1.13	1.22	8.4	1.5	-
19	3.9	6.18	1.18	1.25	10.8	3.6	-
20	3.3	2.81	1.12	1.23	9.2	2.8	-
21	3.4	4.22	1.13	1.15	8.9	1.5	-
22	1.9	1.93	1.09	1.35	5.7	1.3	-
23	2.6	1.91	1.11	1.59	9.0	2.2	-
24	2.2	3.49	1.11	1.18	5.1	0.5	-
25	2.0	3.52	1.16	1.30	5.3	0.6	-
26	1.7	3.91	1.12	1.18	5.7	0.5	-
27	1.2	2.42	1.09	1.15	3.0	0.4	-
28	4.0	3.43	1.12	1.30	22.7	7.2	melanoma
29	1.1	2.77	1.08	1.10	3.5	0.3	-
30	1.8	3.43	1.16	1.26	5.1	0.7	-
31	1.5	2.63	1.09	1.18	3.4	0.6	-
32	2.3	1.75	1.06	1.23	5.6	1.3	melanoma
33	2.4	2.05	1.08	1.21	8.4	2.4	-
34	2.9	3.12	1.14	1.24	8.3	3.0	-
35	1.6	2.60	1.11	1.18	5.5	0.5	-
36	1.6	2.39	1.17	1.45	5.4	0.9	-
37	1.1	2.53	1.15	1.32	3.5	0.6	-
38	4.0	5.39	1.12	1.24	18.9	4.1	-
39	2.6	2.68	1.14	1.30	8.3	1.1	-
40	1.3	2.61	1.15	1.29	3.1	0.5	-




Figure 5.6. The most significant indentation/protrusion segments. The tested lesion borders are presented in Fig. 5.2. Their corresponding MSII's are listed in Table 5.2.

Table 5.3. Spearman coefficients for all test data

	Average clinical evaluation	CI	FD	SFD	OII
CI	0.50				
FD	0.21	0.60			
SFD	0.03	-0.13	0.58		
OII	0.88	0.38	0.17	0.10	
MSII	0.82	0.23	0.18	0.22	0.88




Figure 5.7. Plotting the tested measurements vs. clinical evaluation. The CI, FD, SFD, OII, MSII are plotted (in the x-axis) against the average clinical evaluation index (in the y-axis) in subfigures (a) to (e) separately.

Table 5.4. Multiple linear regression analysis results. Details of the analysis results are discussed in Section 5.3.4.

Dependent variable : average clinical evaluation				
N : 36				
Independent variable	Estimate (SE)	<i>p</i>	Excluded variables	<i>p</i>
OII	0.315 (0.032)	0.000	MSII	0.779
			CI	0.066
			FD	0.973
			SFD	0.171

In order to achieve a reliable gold standard for our user study, we asked dermatologists to classify a lesion border outline into a small score scale, ranging from 1 to 4. Other than the most benign looking contour group (score value 1) and the most malignant looking contour group (score value 4), there were only two intermediate groups. Dermatologists could confidently assign a score value to a tested lesion border based on their own subjective cut points between groups. Judging from the clinical evaluation results in Table 5.1, this score scale worked well and the twenty dermatologists achieved a good agreement. Among the 40 lesion borders, 36 of them had a majority agreement, i.e. agreed by at least 11 dermatologists. More importantly, 39 lesion borders had a score difference of at most 2. This implied that the dermatologists' cut points were *close* to each other. The standard deviation for the clinical evaluation of the 40 lesion borders ranged from 0.00 to 0.73, with the average equaled to 0.50. Furthermore, the high Kendall *W* statistic, $W = 0.77$ ($p = 0.000$), confirmed the good agreement among the clinical evaluations. With the assurance of a good agreement, the gold standard was set up by averaging the clinical evaluations for each lesion border.

5.3.2. Relationship between automatic methods and the average clinical evaluations

In Fig. 5.7, the plots of the tested measurements against the average clinical evaluation showed that no tested measurement achieved a perfect correlation with the gold standard. However, the OII and the MSII had a better linear relationship with the average clinical evaluation than the

other three tested measurements. Similarly, from the Spearman coefficients listed in Table 5.3, the OII had the highest correlation coefficient, 0.88, with the clinical evaluation among all tested measurements. The closest second was the MSII, 0.82.

The CI, FD, and SFD achieved much lower correlation coefficients, 0.50, 0.21, and 0.03, respectively, with the clinical evaluation. These three measures had problems in detecting large structure indentations and protrusions. For example, they failed to properly measure the prominent structure irregularities in tested border 28 of Fig. 5.2, where from Table 5.2, the CI = 3.43, FD = 1.12, SFD = 1.30. These three tested measurements were very close in value to those for border 40 (CI = 2.61, FD = 1.15, SFD = 1.29), which had no structure irregularity, but only texture irregularities. Furthermore, it was surprising to discover that the SFD did not perform better than the FD. The SFD's of lesion border 5 (1.24), border 6 (1.26), border 15 (1.40), border 17 (1.35), border 37 (1.32) and border 40 (1.29) were too high for the benign looking borders in the test set. The problem was caused by applying a wrong model to a border that had only small texture variations. This type of border should have only one overall fractal dimension. When the data points of the log-log plot were separated into two groups, there were too few points to accurately estimate the SFD. Removing these six lesion borders, the Spearman's coefficient for the SFD and the average clinical evaluation improved to 0.30.

5.3.3. A single point estimation vs. a measurement set

The CI and the FD are well-known shape descriptors for the overall border roughness. Both methods compute a single point estimation without actually identifying the indentations and protrusions on the lesion border. A higher value implies a rougher border with the existence of indentations and protrusions. However, a single point estimation can be easily skewed if the variance of the border ruggedness is large. For example, a lesion with a large protrusion on a relatively smooth border, such as lesion border 20, might have the single point estimation dampened by the smoothed portion. The low CI (2.81) and FD (1.12) values were mainly caused by the smooth border that had few texture irregularities.

On the other hand, the new irregularity index method identifies all indentations and protrusions on the lesion border. Because each irregularity is carefully analyzed independently, this method not only returns a more accurate overall measurement, which represents the roughness for the entire border, but also gives a set of measurements for all global irregularities (see Fig. 4.14) that fully describes the complexity of the lesion border. In particular, combining the OII and the MSII may deduce many interesting properties of the lesion borders. For example, the OII (5.5) and the MSII (0.5) of lesion border 35 indicated that it had no structure irregularity, but a lot of texture irregularities. A similar situation holds for lesion border 2 (OII=5.5, MSII=0.6). However, lesion border 32 with the OII = 5.6 and the MSII = 1.3 implied that the lesion border had a larger irregularity, but less texture irregularities than lesion borders 35 and 2. On the other hand, the OII = 22.7 and the MSII = 7.2 for lesion border 28 suggested the border had some major indentations and protrusions. Furthermore, the rich set of measurements for the global irregularities can be used to infer other border properties such as enumerating the number of *large* or *medium* irregularities.

5.3.4. Linear regression model

5.3.4.1. Single group analysis

To determine which tested measurement or a linear combination of tested measurements best predicted the average clinical evaluations, a linear regression model for the average clinical evaluations was computed. However, when all lesion borders ($N = 40$) were used in the analysis, the assumptions for linear regression were violated. For example, plots of the OII and the MSII against clinical evaluation in Fig 5.7d and 5.7e clearly showed that there was no linear relationship between the average clinical evaluation and the OII and the MSII. The problem was caused by limiting the evaluation scores to 4. Limiting the evaluation scores helped the dermatologists properly evaluate the lesion border; however, the side effect was an artificial capping on the highly irregular borders to the score of 4, which might not accurately reflect the degree of the irregularity. For example, the top four data points in Fig. 5.7d and Fig. 5.7e were

lesion borders 4, 19, 28 and 38, which had highly irregular contours. If there were no capping in the clinical evaluation, their actual evaluation scores might exceed 4. The capping produced a flattening effect on Fig. 5.7d and Fig. 5.7e. In order to analyse the data set properly, the lesion borders should be divided into two groups. The first group consisted of the lesion borders 4, 19, 28 and 38 and the second group consisted of the rest of 36 lesion borders.

5.3.4.2. Subgroup analysis

It is difficult to generate any statistically significant results with only 4 data points. Therefore, instead of running a regression model, the visual inspection method was used with the first group of 4 lesion borders. Fig. 5.7 was re-plotted in Fig. 5.8 with the first data group denoted by O. All these data points were at the top of each subfigure because their average clinical evaluations were close to 4. For the subfigure of CI, Fig. 5.8a, there was no obvious relationship between the points marked as O's and the average clinical evaluations because the range of the 4 data points were large. Even though the range of the 4 data points for the FD⁴² and the SFD were relatively small, these data points were located in the middle of the entire FD and SFD ranges with respect to the entire data set. On the other hand, the OII and the MSII showed a clear relationship as these four data points were the highest OII and MSII values among all lesion borders.

The second data group (N = 36) was analysed by a multiple linear regression analysis. One step was required to complete the analysis and the selected independent variable was OII. The parameters for OII and the excluded variables are reported in Table 5.4. The regression line is plotted in Fig. 5.9. A second attempt has been made to fit the linear regression model after also excluding the lesion borders 5, 6, 15, 17, 37, and 40, where the SFD method had problems in generating proper indices. However, no significant changes could be identified in the second regression results. The OII was selected again as the sole independent variable with

⁴² Two of the data points for FD were almost overlapped; therefore, only three O's were clearly recognized.

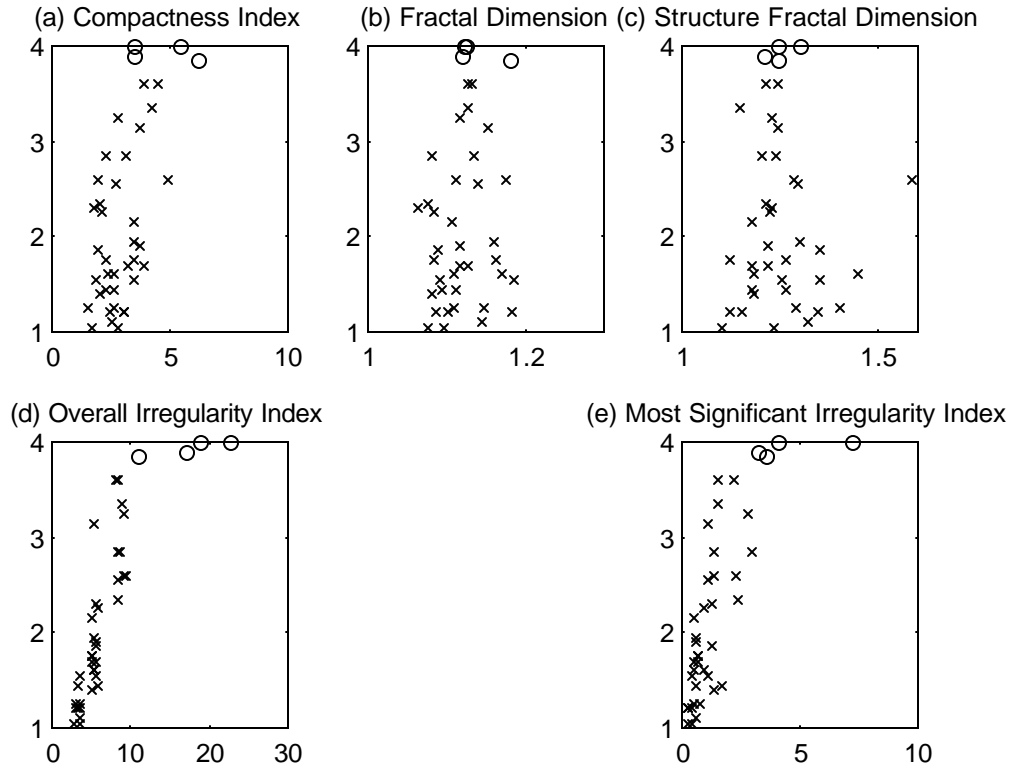


Figure 5.8. The re-plot of Fig. 5.7. The top 4 data points, the first data group in the linear regression analysis, is plotted as O. These points had the highest average clinical evaluations. For the fractal dimension subfigure (b), two of the points (marked as O) are almost overlapped; therefore, only three points are clearly shown.

the parameter estimated as 0.316 ($p = 0.000$). According to the linear regression analysis, the OII was the best predictor for the average clinical evaluations.

Fig. 5.9 also reveals an interesting observation among the three groups of data points. The groups marked with ∇ , $+$, and Δ correspond to the low, middle and high OII values, respectively. (See Section 5.1.2 for the data selection scheme.) The OII has a better prediction power when its value is small, i.e. the deviation from the data point to the regression line increases as the OII increases. This phenomenon may be caused by the algorithm design for the OII, which accumulates the individual irregularity index for every indentation and




Figure 5.9. Predicting average clinical evaluation by OII. The regression line from the first regression model was plotted (N = 36). The data points marked by ∇ were a group of lesion borders with the small OII values. The data points marked by + were a group of lesion borders with the middle OII values. The data points marked by Δ were a group of lesion borders with large OII values. See Section 5.1.2 for the data selection scheme.

protrusion. Therefore the error of the OII is proportional to the number of irregularities along the border. Highly irregular borders tend to have a larger error than smooth borders.

5.3.5. Overall irregularity index and histology

Among all tested measurements, the OII achieved the best correlation with our gold standard, the average clinical evaluation. It would be interesting to compare the OII with histological results of the lesions. Unfortunately, because our data collection period for the 40 lesions spanned four and a half years (from February 1994 to November 1998) and the pathological

assessments were performed in 5 institutions (Vancouver General Hospital, Royal Columbia Hospital, St. Paul's Hospital, Lions Gate Hospital, and St. Vincent's Hospital), forming a gold histological standard for various types of nevi and melanoma was not feasible.⁴³ Therefore, we grouped the lesions into only two definite diagnostic groups: melanoma and non-melanoma groups. Within our 40 tested lesions, there were 4 melanomas (corresponding to lesion borders 4, 6, 28 and 32 as shown in Table 5.2) and 36 non-melanomas based on pathological assessments.

To assess the discriminatory power of the OII, we had to classify the lesions into melanoma and non-melanoma groups based on the OII values. The selection of a classification method became critical. A simple objective classification method assumes all melanomas would have most irregular border and all non-melanomas would have most regular border [31]. Therefore, the 4 lesions with the highest OII values were classified as melanomas and all other lesions were classified as non-melanomas. With this classification method for melanomas, there were 2 true positives, 2 false positives, 34 true negatives and 2 false negatives. The sensitivity of detecting melanomas [54] was determined as:

$$sensitivity = \frac{true\ positive}{true\ positive + false\ negative} \times 100\% = 50\%, \quad (5.2)$$

and specificity [54] was computed as:

$$specificity = \frac{true\ negative}{true\ negative + false\ positive} \times 100\% = 94\%. \quad (5.3)$$

⁴³ We had difficulty in quantifying the likelihood of melanoma because the pathological reports contained subjective descriptions (e.g. melanocytic hyperplasia, mild/moderate architectural atypia, and cytologic abnormality). Furthermore, the reports were done by many pathologists. Ideally, we should have at least two pathologists reviewing all specimens at a single session. The pathologists should agree on the final assessment using a standard protocol and quantify the assessments.

The above discriminatory analysis⁴⁴ suggested that border shape information had excellent power to screen out non-melanomas, but moderate power to detect melanomas. The analysis result confirmed our knowledge that melanomas cannot be diagnosed properly by a single clinical feature. In particular, some melanomas may have regular borders and some non-melanomas may have irregular borders. Based on a sole clinical feature, there is a high probability of misclassification. To properly diagnose melanomas, we have to incorporate OII with other clinical features (discussed in Sect. 1.3) and subsurface features (discussed in Sect. 1.4.1) in a classifier.

In order to verify the actual performance of the OII, we also need to compare the average clinical evaluations and the histological results. Once again, the simple objective classification method was applied. The average clinical evaluations were ranked and the 4 lesions with the highest scores were classified as melanomas and the rest as non-melanomas. With this classification method, we found that the 40 lesions were classified the same as the OII classification method, i.e., the same 4 lesions were classified as melanomas and the same 36 lesions were classified as non-melanomas. The dermatologists mis-classified the same two pathological melanomas as non-melanomas (false negative) and the same two pathological non-melanomas as melanomas (false positive). Therefore, we confirmed that our algorithm indeed captured the knowledge of expert dermatologists on analysing border shape.

5.3.6. Summary

In spite of its good predictive power, there is a disadvantage for the OII. The algorithm for the OII and the MSII is more complicated than for the CI, FD and SFD. The OII and the MSII are computationally expensive even though the computation time can be reduced by optimization and by using a faster computer language and computer hardware. On the other hand, the CI, the most popular method, has the simplest computation complexity. The CI

⁴⁴ The specificity could be improved to 97% by classifying the 3 lesions with the highest OII scores as melanomas.

achieved a moderate Spearman coefficient (0.50) with the average clinical evaluations. The FD is based on a sound mathematical theory, but it is less sensitive for structural irregularities. The SFD is supposed to be able to detect structure features by analysing the coarse scale data of the log-log plot, but the difficulty of properly determining the cutoff point between the fine and coarse scale data for all lesion borders limits the SFD's performance.

The user study reveals many advantages of the irregularity index method. The OII and the MSII correlate well with experienced dermatologists' evaluations. The Spearman's coefficients are 0.88 and 0.83 for the OII and the MSII, respectively. Furthermore, the OII is the best predictor among all tested measures. Because the OII is defined as the sum of a set of irregularity indices for all global irregularity segments, the OII avoids the *skewing (averaging)* effects of many single point estimators. The single point estimation is dampened by the smoothed portion of the lesion border. Another advantage of the new method is that the algorithm pinpoints and highlights a potential problematic area, such as the most significant irregular segment, and explains the OII value by its individual sub-components, which fully describe the lesion shape. Therefore, physicians can verify the highlighted irregular segments and their indices before making the final diagnosis. The detailed information provided by the new indices may be useful for a computer-aided diagnostic device.

Chapter 6

Conclusion and Future Work

The most effective method to reduce mortality for malignant melanomas is early detection since the survival rate is inversely proportional to the thickness of the lesion. Many physicians have advocated some kind of automatic early diagnostic aided systems to improve the diagnostic accuracy and to combat the rapidly increasing incidence rate. Many experimental classifiers have been attempted [33, 38, 56, 58, 129, 149, 150]. The powers of these systems depend on the input features. For my thesis, I focus on one of the important clinical features, border irregularity, which may suggest the malignancy of the skin lesion.

The tasks of designing and implementing automatic procedures to measure lesion borders can be divided into three parts. First, a set of melanocytic images spanning from benign nevi to malignant melanomas have been collected and automatic programs for hair removal and lesion border extraction have been designed. Second, a new shape measure called Irregularity Index (II) has been developed using an extension to classic curvature scale-space images. This method directly locates all indentation and protrusion segments along the border enabling an area-based index to be computed for each irregular segment. From the rich set of measurements, two new shape indices, the overall irregularity index (OII) and the most significant irregularity index (MSII), were derived. Third, the new indices were compared to dermatologists' evaluations in a user study. The result demonstrated that the OII and the MSII vastly outperform other lesion shape descriptors.

6.1. Contribution and originality

In this section, the contribution and the originality of the project is discussed:

1. We constructed new border shape descriptors for melanocytic lesions that were sensitive to structure and texture irregularities by detecting all local and global indentation/protrusion segments. The new method combined and extended many computer shape analysis techniques:

- Object partitioning approach: We extended the border partition rules proposed by Hoffman and Richards, and Siddiqi and Kimia to address indentation segments because their partition rules included only protrusion segments. With the new partition rule, a lesion border was decomposed into a set of indentation and protrusion segments, which are the important clinical features for diagnosing malignant melanomas.
- Multi-scale approach: We extended the classic curvature scale-space approach, which investigates point features, to capture all global and structural indentation and protrusion curve segments of a rugged border. The extended curvature scale-space image differs significantly from the classic curvature scale-space image because of different investigated features. The extended curvature scale space image allowed us:
 - to identify all indentation and protrusion segments, defined by our partitioning rule at each smoothing step.
 - to track the evolution of all indentation and protrusion segments by linking the segments across the smoothing scales.
 - to map the segments back to their *true* positions at the original border.
 - to place the segments in hierarchical structures that helped understand the lesion shape. The segments at the top of the hierarchical structure are considered as the global segments, while the segments at the bottom of the hierarchical structure are the local segments.
 - to investigate the segment properties at all smoothing scales.
- Area-based approach: The construction of the irregularity index was based on the irregularity area.

By combining the above approaches, a new technique to measure border irregularity was designed. The resultant method possessed many advantages over other shape measures:

- The OII and the MSII had high correlation with experienced dermatologists' evaluation of the lesion border in a user study. Moreover, when comparing to other common shape descriptors, the OII was the best predictor for the clinical evaluation of lesion borders in a multiple linear regression model.

- The multi-scale method handled a rugged border and was sensitive to structure irregularities.
 - The irregularity index method was stable because of the area-based approach. An irregularity index was proportional to its irregularity area of an irregular segment.
 - The OII was linear because it summed all the irregularity areas of a lesion. Furthermore, the summation avoided the *skewing (averaging)* effect of a single point estimator, which could be dampened by the smoothing section of the border.
 - The irregularity method actually detected all local and global indentations and protrusions. This set of irregularity segments provided a rich description of the lesion border shape and could be used to derive other parameters for the border shape.
 - The method offered an extra feature: localization of the significant indentations and protrusions, which might be useful for pinpointing problematic areas of a melanocytic lesion in an expert-system type of diagnostic device. Physicians could verify the highlighted area before the final diagnosis were made.
2. We collected a set of skin image data ranging from clinically benign nevi, pathologically benign nevi, pathologically dysplastic nevi and malignant melanomas for future investigation.
 3. The preprocessor DullRazor was designed to reduce the interference of thick dark hairs from skin images, for subsequent analysis programs. Without DullRazor, hairs have to be physically removed by shaving, which is uncomfortable and time consuming. Also it is not possible to remove hairs from existing images without the software approach.
 4. We constructed a segmentation program for the pigmented lesion images. The program extracted the lesion for further analysis.
 5. Forming a reliable gold standard has been an issue for clinical evaluation of lesion border contours. Claridge et al. [32] reported low agreement among expert clinicians, when they were asked to sort lesion borders contours in the order of increasing border irregularity. Such a low agreement cast doubt on forming a credible gold standard in a clinical evaluation setting. For the validation step in this project, we adopted a different experiment design. Dermatologists were asked to score the lesion border on a 4-point scale. The narrow

score scale gave enough separation to the degrees of irregularity and helped dermatologists provide a reliable evaluation.

6.2. Future work

My thesis is part of a larger project: to develop an automatic non-invasive, in-vivo diagnostic device for malignant melanomas. Work developed so far can be the base of the larger project. Of course, improvements can still be made. Also some of the work can be applied to other areas.

DullRazor was developed to simplify the segmentation task by removing dark thick hairs from skin images. However, thick light-coloured hairs such as the one shown in Fig. 6.1 may interfere with some automatic analyses. Extending DullRazor to such a task could be necessary for this type of skin images.




Figure 6.1. A skin image covered by light coloured hairs

There are other ways to extend DullRazor. It could be used to measure hair growth on a shaved patch of scalp for treatment of alopecia. Present accepted methodology involves either manual counting of hair and/or shaving and weighing the hairs from the scalp. The shaving and weighing procedure is complicated by the requirement of separating the skin cells from the hairs prior to weighing. An automatic procedure using scalp images could simplify this problematic task greatly. Another application for DullRazor might allow for the development of

an index related to the vascularity of images of the retina. This could be useful in monitoring the progression of retinal disease, or the response of such disease to therapeutic intervention.

Automatic segmentation is still a challenging active research topic in medical imaging and computer vision. However, without a perfect automatic segmentation program, many analysis programs cannot have a reliable input.

Instead of measuring the lesion border shape, the irregularity index approach may be able to measure the colour shape inside a melanocytic lesion. Malignant melanomas tend to have multiple colours and each colour could have an irregularity shape. Outlining all internal colour edges of a lesion and computing the corresponding OII for the colour edges may provide a better malignancy indicator for a melanocytic lesion.

Although the irregularity index is designed for melanocytic lesions, the shape measure can also be applied to other medical related problems, such as differentiating the malignancy of other solid tumors. For example, the new algorithm can be used to detect and measure the spikes of a breast mass from a mammogram. Furthermore, since the formation of the methodology depends only on a planar closed curve, it can also be used as a shape descriptor for other general 2-D image analysis problems, such as identifying the largest bay, the most significant indentation, on an aerial map. The methodology can potentially be a powerful tool for many medical and scientific applications.

Appendix A

FORM #5

SIMON FRASER UNIVERSITY

INFORMATION SHEET FOR SUBJECTS

Title of Project: Evaluation of melanocytic lesion images

One clinical feature of superficial spreading melanomas is irregularity of edge, often associated with colour variation. This study is an attempt to assess the usefulness of edge information alone, in the absence of colour changes, as a clinical marker of melanoma probability.

These figures are the border outlines of melanocytic lesions. The contours are extracted by an automatic computer program. Because the size of a lesion is itself a marker of risks, all lesions are standardized and enlarged to the same enclosed area when they are printed on the paper. There are two types of variations along a lesion border: global and fine variations. The fine variations could be artifact that is caused by the automatic segmentation program and/or the enlargement process. Therefore, these fine variations should have a less weight than the global variations during the evaluation.

Please evaluate these border contours, rate them individually in the scale of 1 to 4, and record the rating in the space provided just beneath the border. The lowest value of 1 implies the most benign looking nevus border contour, and the highest scale of 4 implies the most severe case with the highest probability of a melanoma contour. There are 4 pages of border contours. At least one contour on each page is a 1, and at least one is a 4. Please find a 1 and a 4 contour, and record that rating before assessing other contours on that page. Do this for each of the four pages.

There is no side effect or risk for this research procedure. Furthermore, there is no personal benefit for the participants. However, information and knowledge gained from the experiment, in the future, may be beneficial to patients with melanocytic lesions.

SIMON FRASER UNIVERSITY
INFORMED CONSENT BY SUBJECTS TO PARTICIPATE
IN A RESEARCH PROJECT OR EXPERIMENT

The University and those conducting this project subscribe to the ethical conduct of research and to the protection at all times of the interests, comfort, and safety of subjects. This form and the information it contains are given to you for your own protection and full understanding of the procedures. Your signature on this form will signify that you have received an *Information Sheet for Subjects* which describes the procedures, possible risks, and benefits of this research project, that you have received an adequate opportunity to consider the information in the *Information Sheet for Subjects*, and that you voluntarily agree to participate in the project.

Any information that is obtained during this study will be kept confidential to the full extent permitted by law. Knowledge of your identity is not required. You will not be required to write your name or any other identifying information on the research materials. Materials will be held in a secure location.

Having been asked by **Tim Lee of the School of Computing Science, Simon Fraser University** to participate in a research project experiment, I have read the procedures specified in the *Information Sheet for Subjects*.

I understand the procedures to be used in this experiment and there is no personal risks or benefits to me in taking part.

I understand that I may withdraw my participation in this experiment at any time.

I also understand that I may register any complaint I might have about the experiment with the researcher named above or with **Dr. James Delgrande (604 291-4335) Chair of the School of Computing Science, Simon Fraser University**.

I may obtain copies of the results of this study, upon its completion, by contacting Tim Lee of School of Computing Science, Simon Fraser University.

I have been informed that the research material will be held confidential by the Principal Investigator.

I understand that my supervisor or employer may require me to obtain his or her permission prior to my participation in a study such as this.

I agree to participate by evaluating the set of melanocytic lesion images into a scale from 1 to 4. The lowest value of 1 implies the most benign looking nevus, and the highest scale of 4 implies the most severe case with the highest probability of being a melanoma.

NAME (please type or print legibly): _____

ADDRESS: _____




SIGNATURE: _____ **WITNESS:** _____




DATE: _____

Evaluate the following lesion border contours using a scale of 1 to 4

1 : most benign looking nevus border contour

4 : highest probability of a melanoma contour

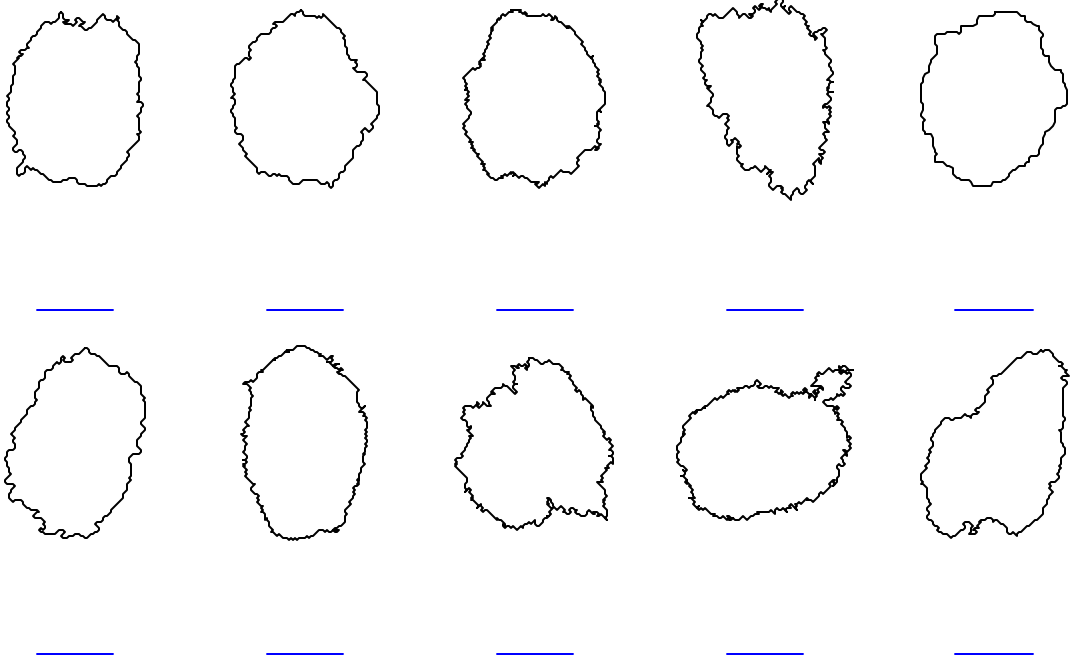




Evaluate the following lesion border contours using a scale of 1 to 4

1 : most benign looking nevus border contour

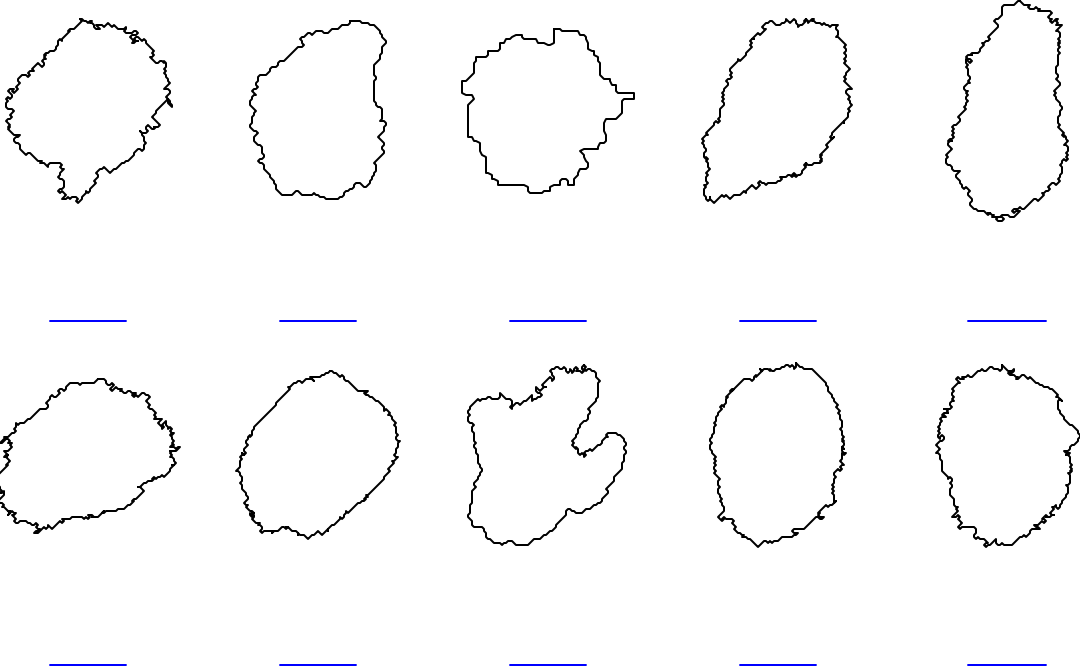
4 : highest probability of a melanoma contour



Evaluate the following lesion border contours using a scale of 1 to 4

1 : most benign looking nevus border contour


4 : highest probability of a melanoma contour



Evaluate the following lesion border contours using a scale of 1 to 4

1 : most benign looking nevus border contour

4 : highest probability of a melanoma contour



References

- [1] P.H. Andersen. Reflectance spectroscopic analysis of selected experimental dermatological models with emphasis on cutaneous vascular reactions. *Skin Research and Technology* **3**(sup 1), 1997, p. 1-58.
- [2] M. Anderson. *Task-oriented lossy compression of magnetic resonance images*. M.Sc. thesis, School of Computing Science, Simon Fraser University, 1995.
- [3] R.R. Anderson and J.A. Parrish. The optics of human skin. *The Journal of Investigative Dermatology* **77**, 1981, p. 13-19.
- [4] L. Andreassi, R. Perotti, M. Burrioni, G. Dell'Eva, and M. Biagioli. Computerized image analysis of pigmented lesions. *Chronica Dermatologica* **5**(1), 1995, p. 11-24.
- [5] G. Arce and R. Foster. Detail-preserving ranked-order based filters for image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing* **37**(1), 1989, p. 83-98.
- [6] G. Argenziano, G. Fabbrocini, P. Carli, V.D. Giorgi, and M. Delfino. Epiluminescence microscopy: Criteria of cutaneous melanoma progression. *Journal of American Academy of Dermatology* **37**, 1997, p. 68-74.
- [7] B.K. Armstrong and D.R. English. Cutaneous malignant melanoma. In *Cancer Epidemiology and Prevention*, D. Schottenfeld and J. Joseph F. Fraumeni, Eds., 2nd ed. New York: Oxford University Press, 1996, p. 1282-1312.
- [8] H. Asada and M. Brady. The curvature primal sketch. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**, 1983, p. 2-14.
- [9] F. Attneave. Some informational aspects of visual perception. *Psychol. Rev.* **61**, 1954, p. 183-193.
- [10] J. Babaud, A.P. Witkin, M. Baudin, and R.O. Duda. Uniqueness of the Gaussian kernel for scale-space filtering. *IEEE Trans. Pattern Analysis and Machine Intelligence* **8**(1), 1986, p. 26-33.
- [11] F.A. Bahmer, P. Fritsch, J. Kreusch, Hubert Pehamberger, C. Rohrer, I. Schindera, J. Smolle, P. Soyer, and W. Stolz. Terminology in surface microscopy. *Journal of American Academy of Dermatology* **23**(6 Part 1), 1990, p. 1159-1162.
- [12] C.M. Balch, A. Houghton, and L. Peters. Cutaneous melanoma. In *Cancer Principle & Practice of Oncology*, V.T.D. Jr., S. Hellman, and S.A. Rosenberg, Eds., 3rd ed: J.B. Lippincott Company, 1989, p. 1499-1542.

- [13] D. Ballard and C. Brown. *Computer Vision*: Prentice-Hall, 1982.
- [14] BC Cancer Agency. Cancer management manual: melanoma, 1999.
<http://www.bccancer.bc.ca/cmm/skin/01.shtml>. (Accessed at July 2000).
- [15] M. Berwick, C.B. Begg, J.A. Fine, G.C. Roush, and R.L. Barnhill. Screening for cutaneous melanoma by skin self-examination. *Journal of National Cancer Institute* **88**(1), 1996, p. 17-23.
- [16] M. Binder, M. Schwarz, A. Winkler, A. Steiner, A. Kaider, K. Wolff, and H. Pehamberger. Epiluminescence microscopy: a useful tool for the diagnosis of pigmented skin lesions for formally trained dermatologists. *Archives Dermatology* **131**, 1995, p. 286-291.
- [17] M. Binder, A. Steiner, M. Schwarz, S. Knollmayer, K. Wolff, and H. Pehamberger. Application of an artificial neural network in epiluminescence microscopy pattern analysis of pigmented skin lesions: a pilot study. *British Journal of Dermatology* **130**, 1994, p. 460-465.
- [18] A. Bittorf, N.C. Krejci-Papa, and T.L. Diepgen. Storage and retrieval of digital images in dermatology. *Skin Research and Technology* **1**, 1995, p. 192-199.
- [19] H. Blum and R.N. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition* **10**, 1978, p. 167-180.
- [20] A. Bono, S. Tomatis, C. Bartoli, G. Tragni, G. Radaelli, A. Maurichi, and R. Marchesini. The ABCD system of melanoma detection: A spectrophotometric analysis of the asymmetry, border, color, and dimension. *Cancer* **85**(1), 1999, p. 72-77.
- [21] A. Bowser. Computerized dermatoscopes could simplify lesion analysis: Developer, however, wondering if devices will catch on in the United States. *Dermatology Times* 1998 October, p. 10-11.
- [22] M. Brady. Criteria for representations of shape. In *Human and Machine Vision*, J. Beck, B. Hope, and A. Rosenfeld, Eds. New York: Academic Press, 1983.
- [23] M. Brady. Representing shape. In *Robotics and Artificial Intelligence*, M. Brady, L.A. Gerhardt, and H.F. Davidsion, Eds. Berlin Heidelberg: Springer-Verlag, 1984.
- [24] N. Brady, F. Bossen, and N. Murphy. Context-based arithmetic encoding of 2D shape sequences. In IEEE International Conference on Image Processing, Santa Barbara, California, 1997, p. 29-32.

- [25] D. Brahmi, C. Serruys, N. Cassoux, A. Giron, R. Triller, P. Lehoang, and B. Fertil. Extraction of features from medical images using a modular neural network approach that relies on learning by sample. In *SPIE Medical Imaging 2000*, San Diego, 2000.
- [26] C.B. Caldwell, S.J. Stapleton, D.W. Holdsworth, R.A. Jong, W.J. Weiser, G. Cooke, and M.J. Yaffe. Characterisation of mammographic parenchymal pattern by fractal dimension. *Phys. Med. Biol* **35**(2), 1990, p. 235-247.
- [27] N. Casinelli, M. Ferrario, T. Tonelli, and E. Leo. A possible new tool for clinical diagnosis of melanoma: the computer. *Clinical Diagnosis of Melanoma* **16**(2), 1987, p. 361-367.
- [28] B.B. Chaudhuri and N. Sarkar. Texture segmentation using fractal dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**(1), 1995, p. 72-77.
- [29] C.C. Chen. Improved moment invariants for shape discrimination. *Pattern Recognition* **26**(5), 1993, p. 683-686.
- [30] M.H. Chen and P.F. Yan. A multiscale approach based on morphological filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**(7), 1989, p. 694-700.
- [31] E. Claridge, P.N. Hall, M. Keefe, and J.P. Allen. Shape analysis for classification of malignant melanoma. *Journal Biomed. Eng.* **14**, 1992, p. 229-234.
- [32] E. Claridge, J.D.M. Smith, and P.N. Hall. Evaluation of border irregularity in pigmented skin lesions against a consensus of expert clinicians. In *Proceedings of Medical Image Understanding and Analysis*, Leeds, U.K., 1998, p. 85-88.
- [33] O. Colot, R. Devinoy, A. Sombo, and D. Brucq. A colour image processing method for melanoma detection. In *The First International Medical Imaging Computing and Computer-Assisted Intervention*, Cambridge, MA, 1998, p. 562-569.
- [34] J.B. Dawson, D.J. Barker, D.J. Ellis, E. Grassam, J.A. Cotterill, G.W. Fisher, and J.W. Feather. A theoretical and experimental study of light absorption and scattering by in vivo skin. *Phys. Med. Biol.* **25**(4), 1980, p. 695-709.
- [35] A.P. Dhawan. An expert system for the early detection of melanoma using knowledge-based image analysis. *Analytical and Quantitative Cytology and Histology* **10**(6), 1988, p. 405-416.
- [36] W. Dummer, H.J. Blaheta, B.C. Bastian, Tamara Schenk, E.B. Brocker, and W. Remy. Preoperative characterization of pigmented skin lesions by epiluminescence microscopy and high-frequency ultrasound. *Archives Dermatology* **131**, 1995, p. 279-285.

- [37] D.L. Ellsworth and T.A. Manolio. The emerging importance of genetics in epidemiologic research. I. Basic concepts in human genetics and laboratory technology. *Annals of Epidemiology* **9**(1), 1999, p. 1-16.
- [38] F. Ercal, A. Chawla, W.V. Stoecker, H.C. Lee, and R.H. Moss. Neural network diagnosis of malignant melanoma from color images. *IEEE Transactions on Biomedical Engineering* **41**(9), 1994, p. 837-845.
- [39] M.A. Fischler and H.C. Wolf. Locating Perceptually Salient Points on Planar Curves. *IEEE transactions on Pattern Analysis and Machine Intelligence* **16**(2), 1994, p. 113-129.
- [40] M.G. Fleming. Image analysis for melanoma diagnosis. *The Melanoma Letter* **18**(1), 2000, p. 1-3.
- [41] M.G. Fleming, C. Steger, A.B. Cognetta, and J. Zhang. Analysis of the network pattern in dermatoscopic images. *Skin Research and Technology* **5**, 1999, p. 42-48.
- [42] H. Freeman. On the encoding of arbitrary geometric configuration. *IRE Transactions on Electronic Computers* **10**(2), 1961, p. 260-268.
- [43] H. Freeman. Computer processing of line-drawing images. *ACM Computing Surveys* **6**, 1974, p. 57-97.
- [44] R.J. Friedman, D.S. Rigel, and A.W. Kopf. Early detection of malignant melanoma: the role of physician examination and self-examination of the skin. *CA-A Cancer Journal for Clinicians* **35**(3), 1985, p. 130-151.
- [45] R.P. Gallagher, B. Ma, D.I. McLean, P. Yang, V. Ho, J.A. Carruthers, and L.M. Warshawski. Trends in basal cell carcinoma, squamous cell carcinoma, and melanoma of the skin from 1973 through 1987. *Journal of the American Academy of Dermatology* **23**(3), 1990, p. 413-421.
- [46] R.P. Gallagher and D.I. McLean. The epidemiology of acquired melanocytic nevi- a brief review. *Dermatologic Clinics* **13**(3), 1995, p. 595-603.
- [47] R.P. Gallagher, J.K. Rivers, T. Lee, C.D. Bajdik, and D.I. McLean. Broad-spectrum sunscreen use and nevus development in school children: results of a randomized trial. In 13th International Congress on Photobiology and 28th Annual Meeting American Society for Photobiology, San Francisco, 2000, p. 214.
- [48] R.P. Gallagher, J.K. Rivers, T.K. Lee, C.D. Bajdik, D.I. McLean, and A.J. Coldman. Broad-spectrum sunscreen use and the development of new nevi in white children: a randomized controlled trial. *The Journal of the American Medical Association* **283**(22), 2000, p. 2955-2960.

- [49] D.J. Gawkrödger. *Dermatology: An Illustrated Colour Text*, 2nd ed. Hong Kong: Churchill Livingstone, 1997.
- [50] M. Gniadecka. Potential for high-frequency ultrasonography, nuclear magnetic resonance, and Raman spectroscopy for skin studies. *Skin Research and Technology* **3**, 1997, p. 139-146.
- [51] J.E. Golston, W.V. Stoecker, R.H. Moss, and I.P.S. Dhillon. Automatic detection of irregular borders in melanoma and other skin tumors. *Computerized Medical Imaging and Graphics* **16**(3), 1992, p. 199-203.
- [52] A.L. Graps. An Introduction to Wavelets. IEEE Computational Sciences and Engineering 1995, p. 50-61.
- [53] H.P. Greenspan and D.J. Benney. *Calculus - An Introduction to Applied Mathematics*: McGraw Hill Inc., 1973.
- [54] C.M. Grin, A. Kopf, B. Welkovich, R. Bart, and M. Levenstein. Accuracy in the clinical diagnosis of malignant melanoma. *Archives Dermatology* **126**, 1990, p. 763-766.
- [55] J.F. Guilloid, A.M. Skaria, D. Salomon, and J.H. Saurat. Epiluminescence videomicroscopy: black dots and brown globules revisited by stripping the stratum corneum. *Journal of the American Academy of Dermatology* **36**(3), 1997, p. 371-377.
- [56] D. Guthowicz-Krusin, M. Elbaum, P. Szwaykowski, and A.W. Kopf. Can early malignant melanoma be differentiated from atypical melanocytic nevus by in vivo techniques? Part II. Automatic machine vision classification. *Skin Research and Technology* **3**, 1997, p. 15-22.
- [57] P. Hall. Clinical diagnosis of melanoma. In *Diagnosis and Management of Melanoma in Clinical Practice*. London and New York: Springer-Verlag, 1992, p. 35-52.
- [58] P.N. Hall, E. Claridge, and J.D.M. Smith. Computer screening for early detection of melanoma-is there a future? *British Journal of Dermatology* **132**, 1995, p. 325-338.
- [59] A.C. Halpern. Imaging analysis: marriage between the computer and the dermoscope. *The Melanoma Letter* **18**(1), 2000, p. 1.
- [60] P.F. Hanrahan, P. Hersey, S.W. Menzies, A.B. Watson, and C.A. D'Este. Examination of the ability of people to identify early changes of melanoma in computer-altered pigmented skin lesions. *Archive Dermatology* **133**, 1997, p. 301-311.

- [61] R.M. Haralick and L.G. Shapiro. *Computer and Robot Vision*, vol. I: Addison-Wesley,, 1992.
- [62] M.F. Healsmith, J.F. Bourke, J.E. Osborne, and R.A.C. Graham-Brown. An evaluation of the revised seven-point checklist for the early diagnosis of cutaneous malignant melanoma. *British Journal of Dermatology* **130**, 1994, p. 48-50.
- [63] D.D. Hoffman and W.A. Richards. Parts of recognition. *Cognition* **18**, 1985, p. 65-96.
- [64] P.L. Horn-Ross, E.A. Holly, S.R. Brown, and D.A. Aston. Temporal trends in the incidence of cutaneous malignant melanoma among Caucasians in the San Francisco-Oakland MSA. *Cancer Causes and Control* **2**, 1991, p. 299-305.
- [65] M.K. Hu. Visual pattern recognition by moment invariants. *IRE Trans. Inform. Theory* **8**, 1962, p. 179-187.
- [66] P. Jackway. *Morphological scale-space with application to three-dimensional object recognition*. Ph.D. thesis, Queensland University of Technology, 1995.
- [67] Steven Jacques. Video imaging with polarized light finds skin cancer margins not visible to dermatologists, 1998. <http://omlc.ogi.edu/omlc/news/feb98/polarization/index.html>. (Accessed at July 2000).
- [68] A. Jain. *Fundamentals of Digital Image Processing*: Prentice-Hall, 1989.
- [69] G. Kaiser. *A Friendly Guide to Wavelets*. Boston: Birkhauser, 1994.
- [70] T. Karamfilov, S. Weichold, K. Karte, W. Visler, and U. Wilina. Remittance spectroscopy mapping of human skin in vivo. *Skin Research and Technology* **5**, 1999, p. 49-52.
- [71] K. Kaspar, M. Vogt, H. Ermert, P. Altmeyer, and S.E. Gammal. Does 100 Mhz sonography allow a more detailed visualization of malignant melanomas? *Skin Research and Technology* **5**, 1999, p. 141.
- [72] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: active contour models. In *Proceedings of the first ICCV*, 1987, p. 259-268.
- [73] R.D. Kenet. Digital imaging in dermatology. *Clinics in Dermatology* **13**, 1995, p. 381-392.
- [74] R.O. Kenet, S. Kang, K. Barney J, T.B. Fitzpatrick, A.J. Sober, and R.L. Barnhill. Clinical Diagnosis of pigmented lesions using digital epiluminescence microscopy. *Archives Dermatology* **129**, 1993, p. 157-174.

- [75] B. Kimia, A. Tannenbaum, and S. Zucker. Shapes, shocks, and deformations, I: The components of shape and the reaction-diffusion space. *International Journal of Computer Vision* **15**, 1995, p. 189-224.
- [76] B.B. Kimia, A.R. Tannenbaum, and S.W. Zucker. Nonlinear shape approximation via the entropy scale space. In Proceedings of the SPIE's Geometric Methods in Computer Vision II, 1993, p. 218-233.
- [77] A. Kjoelen, M.J. Thompson, S.E. Umbaugh, R.H. Moss, and W.V. Stoecker. Performance of AI methods in detecting melanoma. *IEEE Engineering in Medicine and Biology* 1995 July/August, p. 411-416.
- [78] S. Ko and Y. Lee. Center weighted median filters and their applications to image enhancement. *IEEE Transactions on Circuits and Systems* **38**(9), 1991, p. 984-993.
- [79] J.J. Koenderink. The structure of images. *Biological Cybernetics* **50**, 1984, p. 363-370.
- [80] H.K. Koh, L.A. Norton, A.C. Geller, T. Sun, D.S. Rigel, D.R. Miller, R.G. Sikes, K. Vigeland, E.U. Bachenberg, M. P.A, F.S. Billon, G. Goldberg, D.A. Scarborough, W.M. Ramsdell, V.A. Muscarella, and R.A. Lew. Evaluation of the American Academy of Dermatology's national skin cancer early detection and screening program. *Journal of the American Academy of Dermatology* **34**(6), 1996, p. 971-978.
- [81] A.W. Kopf, M. Elbaum, and N. Provost. Editorial: the use of dermoscopy and digital imaging in the diagnosis of cutaneous malignant melanoma. *Skin Research and Technology* **3**, 1997, p. 1-7.
- [82] L. Lam, S.W. Lee, and C.Y. Suen. Thinning methodologies-a comprehensive survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(9), 1992, p. 869-885.
- [83] T. Lee, S. Atkins, R. Gallagher, C. MacAulay, A. Coldman, and D. McLean. Describing the structural shape of melanocytic lesions. In SPIE Medical Imaging 1999, San Diego, 1999, p. 1170-1179.
- [84] T. Lee, S. Atkins, and D. McLean. Measuring structural border irregularity for pigmented skin lesions (Abstract). *Skin Research and Technology* **5**, 1999, p. 143.
- [85] T. Lee, V. Ng, R. Gallagher, A. Coldman, and D. McLean. DullRazor: A software approach to hair removal from images. *Computer Biology and Medicine* **27**, 1997, p. 533-543.

- [86] T. Lee, V. Ng, R. Gallagher, A. Coldman, and D. McLean. DullRazor: removal of dark think hairs from skin images (Abstract). *Skin Research and Technology* **3**, 1997, p. 199.
- [87] T. Lee, V. Ng, D. MaLean, A. Coldman, R. Gallagher, and J. Sale. A multi-stage segmentation method for images of skin lesions. In *IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, Victoria, BC, Canada, 1995, p. 602-605.
- [88] T.K. Lee and M.S. Atkins. A new approach to measure border irregularity for melanocytic lesions. In *SPIE Medical Imaging 2000*, San Diego, 2000, p. 668-675.
- [89] T.K. Lee and M.S. Atkins. A new shape measure for melanocytic lesions. In *Medical Image Understanding and Analysis 2000*, London, England, 2000, p. 25-28.
- [90] F. Leymarie and M.D. Levine. Simulating the grassfire transform using an active contour model. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(1), 1992, p. 56-75.
- [91] M. Leyton. A Process-Grammar for Shape. *Artificial Intelligence* **34**, 1988, p. 213-247.
- [92] M. Leyton. Inferring causal history from shape. *Cognitive Science* **13**, 1989, p. 357-387.
- [93] Z.N. Li, O.R. Zaiane, and Z. Tauber. Illumination invariance and object model in content-based image and video retrieval. *Journal of Visual Communication and Image Representation* **10**, 1999, p. 219-244.
- [94] T. Lindeberg. *Scale-space Theory in Computer Vision*. Boston: Kluwer Academic Publishers, 1994.
- [95] S. Loncaric. A survey of shape analysis technique. *Pattern Recognition* **31**(8), 1998, p. 983-1001.
- [96] C. MacAulay. *Development, implementation and evaluation of segmentation algorithms for the automatic classification of cervical cells*. Ph.D. thesis, Department of Physics, The University of British Columbia, 1989.
- [97] R.M. MacKie. Melanoma Prevention and Early Detection. *British Medical Bulletin* **51**(3), 1995, p. 570-583.
- [98] R. MacLennan, A.C. Green, G.R.C. McLeod, and N.G. Martin. Increasing incidence of cutaneous melanoma in Queensland, Australia. *Journal of the Nation Cancer Institute* **84**(18), 1992, p. 1427-1432.

- [99] J.C. Maize and A.B. Ackerman. *Pigmented Lesions of the Skin*. Philadelphia: Lea & Febiger, 1987.
- [100] B.B. Mandelbrot. *The Fractal Geometry of Nature*. New York: W.H. Freeman and Co., 1982.
- [101] D.I. McLean, R.A. Lew, A.J. Sober, J. Martin C. Mihm, and T.B. Fitzpatrick. On the prognostic importance of white depressed areas in the primary lesion of superficial spreading melanoma. *Cancer* **43**, 1979, p. 157-161.
- [102] F. Mokhtarian. Silhouette-based isolated object recognition through curvature scale space. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**(5), 1995, p. 539-544.
- [103] F. Mokhtarian and A. Mackworth. Scale-based description and recognition of planar curves and two-dimensional shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**(1), 1986, p. 34-43.
- [104] F. Mokhtarian and A.K. Mackworth. A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(8), 1992, p. 789-805.
- [105] F. Mokhtarian and R. Suomela. Robust image corner detection through curvature scale space. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **20**, 1998, p. 1376-1381.
- [106] M. Moncrieff, S. Cotton, P. Hall, and E. Claridge. Extracting the diagnostic features of skin surface microscopy for pigmented skin lesions using spectrophotometric intracutaneous analysis. In *Medical Image Understanding and Analysis*, London, 2000, p. 49-52.
- [107] National Cancer Institute of Canada. *Canadian Cancer Statistics 2000, Toronto, Canada*: National Cancer Institute of Canada, 2000.
- [108] V. Ng and T. Lee. Measuring border irregularities of skin lesions using fractal dimensions. In *SPIE Photonics China, Electronic Imaging and Multimedia Systems*, Beijing, China, 1996, p. 64-72.
- [109] M.J. Norusis. *Introductory Statistics Guide for SPSS-X Release 3*: SPSS Inc., 1988.
- [110] R.L. Ogniewicz and O. Kubler. Hierarchic Voronoi skeletons. *Pattern Recognition* **28**(3), 1995, p. 343-359.

- [111] J. Ostermann, E.S. Jang, J.S. Shin, and T. Chen. Coding the arbitrarily shaped video objects in MPEG-4. In *IEEE International Conference on Image Processing*, Santa Barbara, California, 1997, p. 496-499.
- [112] D.M. Parkin, S.L. Whelan, J. Ferlay, L. Raymond, and J. Young, editors. *Cancer Incidence in Five Continents*, vol. VII. Lyon: International Agency for Research on Cancer, 1997.
- [113] T. Pavlidis. A review of algorithms for shape analysis. *Computer Graphics and Image Processing* **7**, 1978, p. 243-258.
- [114] T. Pavlidis. Algorithm for shape analysis of contours and waveforms. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **2**(4), 1980, p. 301-312.
- [115] H. Pehamberger, M. Binder, S. Knollmayer, and K. Wolff. Immediate effects of a public education campaign on prognostic features of melanoma. *Journal of the American Academy of Dermatology* **29**(1), 1993, p. 106-109.
- [116] H. Pehamberger, M. Binder, A. Steiner, and K. Wolff. In vivo epiluminescence microscopy: improvement of early diagnosis of melanoma. *The Journal of Investigative Dermatology* **100**(3), 1993, p. 356S-362S.
- [117] H.O. Peitgen, H. Jurgens, and D. Saupe. *Chaos and Fractals*. New York: Springer-Verlag, 1992.
- [118] A.P. Pentland. Fractal-based description of natural scenes. *IEEE Transactions of Pattern Analysis and Machine Intelligence* **6**(6), 1984, p. 661-674.
- [119] C.E. Priebe, J.L. Solka, R.A. Lorey, G.W. Rogers, W.L. Poston, M. Kallergi, W. Qian, L.P. Clarke, and P.R. Clark. The application of fractal analysis to mammographic tissue classification. *Cancer letters* **77**, 1994, p. 183-189.
- [120] D. Puppin, D. Salomon, and J.-H. Saurat. Amplified surface microscopy. *Journal of American Academy of Dermatology* **28**(6), 1993, p. 923-927.
- [121] W. Qian, L. Clarke, M. Kallergi, and R. Clark. Tree-structured nonlinear filters in digital mammography. *IEEE Transactions on Medical Imaging* **13**(1), 1994, p. 25-36.
- [122] B.K. Rao, A.A. Marghoob, W. Stolz, A.W. Kopf, J. Slade, Q. Wasti, S.P. Schoenbach, M. De-David, and R.S. Bart. Can early malignant melanoma be differentiated from atypical melanocytic nevi by in vivo techniques? Part I. Clinical and dermoscopic characteristics. *Skin Research and Technology* **3**, 1997, p. 8-14.

- [123] W. Richards and D.D. Hoffman. Condon Constraints on Closed 2D Shapes. *Computer Vision, Graphics, and Image Processing* **31**, 1985, p. 265-281.
- [124] J.K. Rivers. Melanoma. *Lancet* **347**, 1996, p. 803-807.
- [125] J.K. Rivers and R.P. Gallagher. Public education projects in skin cancer. Experience of the Canadian Dermatology Association. *Cancer* **75**(sup 2), 1995, p. 661-666.
- [126] A. Rosenfeld and A. Kak. *Digital Picture Processing*, vol. 1: Academic Press, 1982.
- [127] E. Ruoslahti. How Cancer Spreads. *Scientific American* 1996 September, p. 72-77.
- [128] T. Sakamaki, S. Hayashi, and R. Hayakawa. Melanin distribution in corneocytes and fine patterns of pigmentation on pigmented areas in pigmentary disorders. *Skin Research and Technology* **5**, 1999, p. 37-41.
- [129] T. Schindewolf, W. Stolz, R. Albert, W. Abmayr, and H. Harms. Classification of melanocytic lesions with color and texture analysis using digital image processing. *Analytical and Quantitative Cytology and Histology* **15**(1), 1993, p. 1-11.
- [130] S. Seidenair, G. Pellacani, and P. Pepe. Digital videomicroscopy improves diagnostic accuracy for melanoma. *Journal of the American Academy of Dermatology* **39**(2, part 1), 1998, p. 175-181.
- [131] S. Seidenari, M. Burrioni, D.E. Giordana, P. Pepe, and B. Belletti. Computerized evaluation of pigmented skin lesion images record by a videomicroscope: comparison between polarizing mode observation and oil/slide mode observation. *Skin Research and Technology* **1**, 1995, p. 187-191.
- [132] C. Serruys, D. Brahmi, A. Giron, J. Vilain, R. Triller, and B. Fertil. Classification of skin tumors by an adaptive approach grounded on a learning-by-example process: application to the detection of melanomas. *SPIE's International Technical Group Newsletter* 1999 November, p. 1-3.
- [133] S. Shao and R.R. Grams. A proposed computer diagnostic system for malignant melanoma (CDSMM). *Journal of Medical Systems* **18**(2), 1994, p. 85-96.
- [134] K. Siddiqi and B.B. Kimia. Parts of visual form: computational aspects. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**(3), 1995, p. 239-251.
- [135] S. Siegel. *Non-parametric Statistics for the Behavioral Sciences*: McGraw-Hill, 1956.

- [136] W. Slue, A.W. Kopf, and J.K. Rivers. Total body photographs of dysplastic nevi. *Archives Dermatology* **124**, 1988, p. 1239-1243.
- [137] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis and Machine Vision*. London: Chapman & Hall, 1995.
- [138] W.V. Stoecker. *Computer Applications in Dermatology*. New York: Igaku-Shoin, 1993.
- [139] W.V. Stoecker and R.H. Moss. Editorial: digital imaging in dermatology. *Computerized Medical Imaging and Graphics* **16**(3), 1995, p. 145-150.
- [140] W.V. Stoecker, R.H. Moss, F. Ercal, and S.E. Umbaugh. Nondermatoscopic digital imaging of pigmented lesions. *Skin Research and Technology* **1**, 1995, p. 7-16.
- [141] G. Strang. Wavelets and dilation equations: a brief introduction. *Siam Review* **31**, 1989, p. 613-627.
- [142] G. Strang. Wavelets. *American Scientist* 1994 April, p. 250-255.
- [143] M.R. Teague. Image analysis via the general theory of moments. *Journal of Optical Society of America* **70**(8), 1980, p. 920-930.
- [144] G.J. Tortora and N.P. Anagnostakos. *Principles of Anatomy & Physiology*, 2nd ed: Canfield Press, 1978.
- [145] H. Voigt and R. ClaBen. Topodermatographic image analysis for melanoma screening and the quantitative assessment of tumor dimension parameters of the skin. *Cancer* **75**(4), 1995, p. 981-988.
- [146] S.S. Wang, P.C. Chen, and W.G. Lin. Invariant pattern recognition by moment Fourier descriptor. *Pattern Recognition* **27**(12), 1994, p. 1735-1742.
- [147] R.A. Weinberg. How cancer arises. *Scientific American* 1996 September, p. 62-70.
- [148] R.A. Weinberg. *One Renegade Cell*. New York: Basic Books, 1998.
- [149] R. White, D.S. Rigel, and R.J. Friedman. Computer applications in the diagnosis and prognosis of malignant melanoma. *Dermatologic Clinics* **9**(4), 1991, p. 695-702.
- [150] R.G. White, D.A. Perednia, and R.A. Schowengerdt. Automated feature detection in digital images of skin. *Computer Methods and Programs in Biomedicine* **34**, 1991, p. 41-60.
- [151] A.P. Witkin. Scale-space filtering. In Proceedings of the 8th International Joint Conference on Artificial Intelligence, 1983, p. 1019-1022.

- [152] K. Xin, K.B. Lim, and G.S. Hong. A scale-space filtering approach for visual feature extraction. *Pattern Recognition* **28**(8), 1995, p. 1145-1158.
- [153] N. Yamaguchi, T. Ida, and T. Watanabe. A binary shape coding method using Modified MMR. In IEEE International Conference on Image Processing, Santa Barbara, California, 1997, p. 504-507.
- [154] C. Zahn and R. Roskies. Fourier descriptors for plane closed curves. *IEEE Transactions on computers* **C-21**, 1972, p. 269-281.
- [155] H. Zeng. *Human Skin Optical Properties and Autofluorescence Decay Dynamics*. Ph.D. thesis, Department of Physics, The University of British Columbia, 1993.