



Michael Schneeberger, Bakk.techn.

Image Based Analysis and Synthesis of Chronic Wounds and Simulation of Wound Development

MASTER'S THESIS

to achieve the university degree of

Diplom-Ingenieur

Master's degree programme: Telematics

submitted to

Graz University of Technology

Supervisor

Univ.-Prof. Dipl.-Ing. Dr. Horst Bischof
Institute of Computer Graphics and Vision

Advisor

Dipl.-Ing. Dr. Martin Urschler
Ludwig Boltzmann Institute for Clinical Forensic Imaging

AFFIDAVIT

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly indicated all material which has been quoted either literally or by content from the sources used. The text document uploaded to TUGRAZonline is identical to the present master's thesis.

Date

Signature

EIDESSTATTLICHE ERKLÄRUNG

Ich erkläre an Eides statt, dass ich die vorliegende Arbeit selbstständig verfasst, andere als die angegebenen Quellen/Hilfsmittel nicht benutzt, und die den benutzten Quellen wörtlich und inhaltlich entnommenen Stellen als solche kenntlich gemacht habe. Das in TUGRAZonline hochgeladene Textdokument ist mit der vorliegenden Diplomarbeit identisch.

Datum

Unterschrift

Abstract

Chronic wound healing disorders are a serious cost factor in the healthcare system worldwide. Beside professional medical treatment, the active cooperation of the affected patients is essential for an optimal healing process. In order to make the patients aware of their direct influence onto wound healing and to increase the compliance there was the need to develop an image based software tool that is able to simulate possible evolutions of a chronic wound depending on personal life circumstances and make them visually experience-able.

In this environment, the present work focuses on image-based wound analysis and wound synthesis and thus makes the essential contribution to the simulation tool developed in the SimuWound research project. Based on a careful selection and evaluation of suitable descriptors for wound description, wounds were dissected into individual tissue types in order to generate synthetic tissue layers using patch based methods. Deploying derived filters from appropriate masks, new synthetic but still realistic wounds were generated. The development of a reliable detection algorithm for reference markers for scale and colour calibration was also an important contribution. Furthermore, various methods for image-based segmentation of affected body parts, smooth and difficult to define wound borders but also of fistulas within a wound have been developed and successfully tested. Finally, the work outlines a method to generate simulations for wound healing and wound deterioration using the proposed layer concept and a suitable parameter based model.

The developed simulation software has already been successfully tested in the clinical environment. It also has the potential for the use in mobile devices or as a valuable tool in the training of prospective physicians and nursing staff in the clinical and private sectors.

Keywords. wound image analysis, wound synthesis, wound development simulation, image completion, segmentation

Kurzzusammenfassung

Chronische Wundheilungsstörungen sind weltweit ein ernstzunehmender Kostenfaktor im Gesundheitssystem. Für einen optimalen Heilungsverlauf ist neben professioneller medizinischer Behandlung die aktive Mitarbeit der Betroffenen unerlässlich. Um den Patienten ihren direkten Einfluss auf die Wundheilung bewusst zu machen und sie zur Mitarbeit zu motivieren, bestand die Notwendigkeit eine bildbasierte Software zu entwickeln, die im Stande ist, mögliche Verläufe einer chronischen Wunde in Abhängigkeit von persönlichen Lebensumständen zu simulieren und visuell erfahrbar zu machen. Die vorliegende Arbeit hat in diesem Umfeld den Fokus auf bildbasierter Wundanalyse und Wundsynthese und leistet somit den wesentlichen Beitrag für das im Forschungsprojekt SimuWound entwickelte Simulationstool.

Basierend auf einer sorgfältigen Auswahl und Evaluierung geeigneter Deskriptoren zur Wundbeschreibung, konnten Wunden in einzelne Gewebstypen zerlegt werden, um daraus mit patchbasierten image completion Methoden synthetische Gewebslayer zu erzeugen. Mithilfe geeigneter Masken und daraus abgeleiteten Filtern wurden so neue synthetische, aber dennoch realitätsnahe Wunden generiert. Ein wesentlicher Beitrag war auch die Entwicklung eines zuverlässigen Detektionsalgorithmus für Referenzmarker zur Maßstabs- und Farbkalibration. Des Weiteren wurden verschiedene Methoden zur bildbasierten Segmentierung von betroffenen Körperteilen, von weich verlaufenden, schwer abzugrenzenden Wundrändern, aber auch von Fisteln innerhalb einer Wunde entwickelt und erfolgreich getestet. Zum Abschluss skizziert die Arbeit eine Methode um mithilfe des vorgestellten Layerkonzepts und einem geeigneten Parametermodell Simulationen für die Wundheilung und Wundverschlechterung zu erzeugen.

Die entwickelte Simulationssoftware wurde bereits erfolgreich im klinischen Umfeld getestet. Sie hat auch das Potenzial für den Einsatz in mobilen Geräten oder als wertvolles Werkzeug bei der Ausbildung von angehenden Ärztinnen und Ärzten sowie Pflegepersonal im klinischen und niedergelassenen Bereich.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Goals	2
1.3	Scope of Work	3
1.4	Diagnoses	4
1.4.1	Diabetic Foot Syndrome	4
1.4.2	Peripheral Artery Occlusive Disease	5
1.4.3	Venous Leg Ulcer	6
1.5	Outline	7
2	Related Work	9
3	Image Acquisition and Marker Detection	11
3.1	Image Acquisition System	11
3.1.1	Requirements	11
3.1.2	Camera Setup	12
3.1.3	Description of Different Marker Types	14
3.2	Marker Detection	16
3.2.1	Detection Approach	16
3.2.2	Maximally Stable Extremal Regions	19
3.2.3	Performance Evaluation	21

4	Wound Analysis	25
4.1	Wound Description	25
4.2	Dividing a Wound Evolution into Phases	26
4.3	Features for Wound Description	27
4.4	Image Based Descriptors for Wounds	29
4.4.1	Wound Size	29
4.4.2	Texture	29
4.4.2.1	Local Binary Patterns	30
4.4.2.2	Local Variance Measure	31
4.4.2.3	Binary Robust Independent Elementary Features	31
4.4.3	Colour	32
4.5	Automatic Wound Classification	32
4.5.1	Cluster Analysis Based on Wound Descriptors	32
4.5.2	Decision Tree Based on Wound Descriptors	32
4.6	Finding Similar Patches	33
4.6.1	Wound Tissue Patch Database	33
4.6.2	Matching Reference Patches Against a Database	34
4.7	Classification of Wound Regions into Tissue Classes	38
4.7.1	Statistical Region Merging	40
4.7.2	Segmentation of Cavities Inside a Wound	40
4.7.3	Colour Based Segmentation	42
5	Wound Synthesis	45
5.1	Skin Segmentation	45
5.1.1	Model Based Skin Detector	46
5.1.2	Region Growing Based Skin Detector	48
5.1.3	Performance Evaluation	51
5.2	Tissue Layer Concept	53
5.2.1	Layer Stack	53
5.2.2	Tissue Layer Masks	54

5.2.3	Creating a Synthetic Wound	55
5.3	Generating Synthetic Wound Tissue	56
5.4	Completion of Wound Areas with Healthy Skin	56
5.5	DFS Wound Border Treatment	58
5.5.1	Wound Border Region Segmentation Using TVSeg	59
5.5.2	Wound Border Region Segmentation Using SRM	61
5.5.3	Creation of DFS Wound Border	62
5.6	Wound Cavity Synthesis	63
6	Wound Development Simulation	65
6.1	Available Image Data	65
6.2	Wound Development Simulation Chain	65
6.3	Simulation Results	70
7	Conclusion and Outlook	73
7.1	Summary of Contribution	73
7.2	Conclusion, Discussion and Outlook	74
A	List of Publications	77
	Bibliography	78

List of Figures

1.1	Main technical tasks of project	3
1.2	Wound examples of diabetic foot diseases	5
1.3	Wound examples of peripheral artery occlusive diseases	6
1.4	Wound examples of venous leg ulcers	7
3.1	Equipment deployed for wound imaging at the hospital	13
3.2	Image acquisition setup in a hospital	14
3.3	Overview of deployed markers	15
3.4	Intermediate steps of marker detection procedure	18
3.5	State flow diagram of marker detection procedure	20
3.6	Examples of wrongly classified marker types	23
3.7	Example of colour field evaluation chart	23
4.1	Local topology of a wound	26
4.2	Wound images manually grouped in phases	27
4.3	Principle of Local Binary Patterns	30
4.4	Decision tree for classifying wound phase	33
4.5	Patch matching results for short-listed descriptors - example 1	35
4.6	Patch matching results for short-listed descriptors - example 2	37
4.7	Patch matching results for short-listed descriptors - example 3	38
4.8	Patch matching results for combined descriptor	39
4.9	Intermediate steps of cavity detection procedure	41
4.10	Cavity detection results	42

4.11	Distribution of wound tissue training data in RGB space	43
4.12	Results of colour based tissue segmentation using SVM	43
5.1	Intermediate steps of skin detection based on Gaussian mixture model . . .	47
5.2	Skin detection results based on a Gaussian mixture model	48
5.3	Results of Statistical Region Merging	49
5.4	Distribution of skin and background training data in RGB space	50
5.5	Skin detection results of SRM approach using a SVM as classifier	50
5.6	Reference examples for skin segmentation performance rating	52
5.7	Performance comparison of different skin segmentation methods based on colour normalised images	53
5.8	Stack of tissue layers for wound synthesis	54
5.9	Creating a synthetic Wound	56
5.10	Wound image completion technique	57
5.11	Wound image completion results	59
5.12	Wound border region segmentation using TVSeg - example 1	60
5.13	Wound border region segmentation using TVSeg -example 2	61
5.14	Wound border region segmentation using SRM	62
5.15	Synthetically generated cavities	63
6.1	State flow diagram of image preprocessing.	68
6.2	State flow diagram of simulation process.	69
6.3	Healing simulation for DFS	70
6.4	Deterioration simulation for DFS	70
6.5	Healing simulation for UCV	71
6.6	Deterioration simulation for UCV	71

List of Tables

3.1	Camera settings	12
3.2	Image set for detector evaluation	21
3.3	Marker detection rates	22
3.4	Marker detection confusion matrix	22
4.1	Classification of DFS wounds into phases	28
4.2	Classification of PAOD wounds into phases	28
4.3	Classification of UCV wounds into phases	29
4.4	Results of cavity detection tests	41
5.1	Performance comparison of different skin segmentation methods	52
6.1	Available image data for DFS	66
6.2	Available image data for PAOD	66
6.3	Available image data for UCV	67

List of Acronyms

API	Application Programming Interface
BRIEF	Binary Robust Independent Elementary Features
DFS	Diabetic Foot Syndrome
DLT	Direct Linear Transformation
DOG	Difference of Gaussian
GMM	Gaussian Mixture Model
GUI	Graphical User Interface
HSV	Hue-Saturation-Value colour model
ID	Identifier
IDF	International Diabetes Federation
IR	Infra-red
LBP	Local Binary Patterns
LOG	Laplacian of Gaussian
MSER	Maximally Stable Extremal Regions
NIRS	Near Infrared Spectroscopy
PAD	Peripheral Artery Disease
PAOD	Peripheral Artery Occlusive Disease
PCA	Principal Component Analysis
PVD	Peripheral Vascular Disease
RGB	Red-Green-Blue colour model
ROI	Region of Interest
SRM	Statistical Region Merging

SVM	Support Vector Machine
TPS	Thin Plate Spline
TTL	Through the Lens
TVSeg	Interactive Total Variation Based Image Segmentation
UCV	Ulcus Cruris Venosum
XYZ	CIE 1931 XYZ color space

Chapter 1

Introduction

Contents

1.1	Motivation	1
1.2	Goals	2
1.3	Scope of Work	3
1.4	Diagnoses	4
1.5	Outline	7

1.1 Motivation

Chronic wounds are a protracted disease causing substantial health care costs world-wide. Currently, in German-speaking countries alone, several million people suffer from chronic wound healing disturbances. The medical condition is mainly caused by diabetes, peripheral arterial occlusive disease, chronic venous disorder, decubitus wounds or postoperative healing disorders.

Besides professional medical treatment, optimal wound healing presupposes an active involvement of the patient. According to wound experts an active participation may notably decrease the treatment period and thus, directly relieve patient and physician as well as ease the burden on the health budget. However, affected people are often discouraged and dejected due to the long-lasting course of their disease. Despite medical care, many chronic wounds take several months, sometimes even more than a year, to heal completely. Unfortunately patients are often not aware that a therapeutic success strongly depends on their active cooperation and compliance. In order to sharpen the patient's

awareness of their influence onto wound healing there was the need to develop an image analysis and synthesis software tool to simulate possible evolutions of a chronic wound and make them visually experienceable. The simulations should encourage patients to actively participate in therapy but also demonstrate the consequences of non-cooperation in a dissuasive manner.

1.2 Goals

As already pointed out, the aim of the simulation tool is to predict possible wound development scenarios based on the actual state. The developed algorithms deployed in the tool had to consider personal influence factors like for example smoking, diabetes or overweight, allowing the simulation of different scenarios. To increase the effect and impression, the patient should not be confronted with an arbitrary wound simulation, but with the simulation of the development of his own wound on his own body part. Furthermore, to emphasize the urgency, not only healing had to be predicted and visualized but possible deterioration – in case of an untreated wound – as well. Finally, the tool should, inter alia, be ready for the use within doctor-patient conversations and thus be easy and intuitively to use. Therefore, the effort for a physician to create a wound simulation should be limited to the generation of an image depicting the affected body part and a Graphical User Interface (GUI) based annotation including wound contour, diagnosis and personal circumstances.

As the quality of the wound course simulation strongly depends on the available wound data for analysis and training, detractions regarding the simulation accuracy/quality are tolerated. Therefore, deviations from a de facto wound development are not critical as long as the visualisation of the wound evolution looks reasonably realistic. An additional restriction affects the simulation of wound deterioration. As all patients were treated in the hospital, there is usually no training data available for deterioration of a wound. For this reason the deterioration process shall basically be simulated in the same manner like healing but in the opposite direction and temporally much faster. A final point concerns calculation time. There is no need for a real-time application, so processing speed is no issue.

1.3 Scope of Work

The development of the simulation tool was subject to project SimuWound at JOAN-NEUM RESEARCH. Substantial parts of the work are covered in the current thesis. Figure 1.1 provides a retrospective overview of the main technical tasks that had to be solved within the scope of the project. Blue labelled tasks were developed by the author. Yellow tasks were solved within the project team and green labelled assignments were dedicated to other team members than the author.

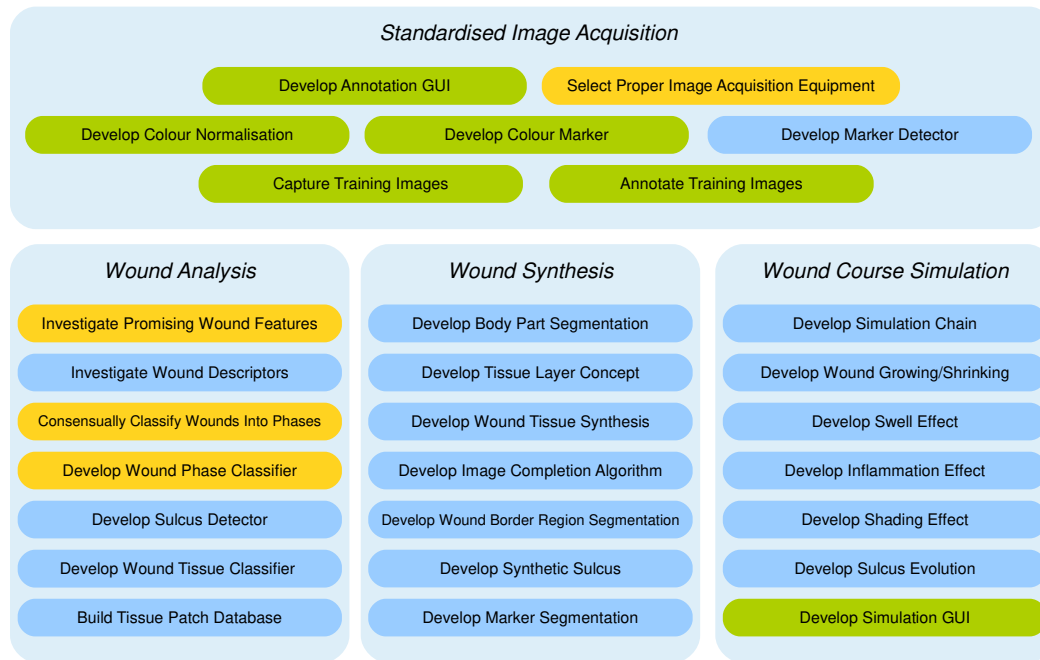


Figure 1.1: Main technical tasks of project. The blue ones are part of this thesis. The yellow ones were solved together within the project team and green ones were dedicated to other team members than the author.

The major contribution of the thesis may be summarized as the analysis and description of visual attributes of a chronic wound with respect to a related diagnosis and the generation of synthetic wounds (based on the extracted features). In addition, the development of an algorithm that alters the wound parameters in order to simulate different wound evolutions depending on given personal circumstances was outlined.

To allow image based analysis and evaluation of different wound evolutions, standardised data acquired in a hospital is made available. These data comprise images of chronic wounds and its borders delineated by a physician. Furthermore, related information about present diagnosis as well as wound course affecting personal circumstances are provided.

In order to reduce the complexity of the assignment, the scope of this work was limited to three diagnoses, and by virtue of being restricted to 2D data, the location of the wound to be examined was restricted to more or less planar areas of the body. So the regions of the body intended for wound simulations comprise – depending on the particular diagnosis – thigh, shank, ankle, sole of the foot and bridge of the foot. The relevant diagnoses are introduced in following section 1.4 and include diabetic foot syndrome, peripheral artery occlusive disease and venous leg ulcer.

Note that some parts of this thesis were already presented in [52]. The focus of that paper was on the newly introduced image completion technique optimised for the simulation of wound development. Since three authors were involved, it has to be mentioned that the assignment of tasks described in Figure 1.1 applies without restriction for the paper as well.

1.4 Diagnoses

In the following, three chronic diseases that have to be considered within the scope of the thesis, are described in short. In order to give the reader a first impression, also related wound images are provided.

1.4.1 Diabetic Foot Syndrome

Diabetes mellitus is a global medical problem. The most recent estimates of the International Diabetes Federation (IDF) in [23] indicate that 8.3% of adults, i.e. 387 million people, suffer from diabetes, and the number of people with the disease is expected to rise beyond 592 million over the next two decades. Furthermore, about 175 million of cases are currently undiagnosed. Accordingly, every second diabetic is unaware of her or his illness and gets into serious health complications without knowing it. One of the most feared secondary diseases of diabetes mellitus is Diabetic Foot Syndrome (DFS), colloquially also referred to as diabetic foot. According to [23] a diabetic foot is defined as "a foot that exhibits any pathology that results directly from diabetes or complication of diabetes". In [24], the International Working Group on the Diabetic Foot defines the disease as the presence of "Infection, ulceration or destruction of deep tissues of the foot associated with neuropathy and/or peripheral arterial disease in the lower extremity of people with diabetes".

The major reasons for diabetic foot syndrome are nerve damages caused by a persistent

high blood glucose level due to untreated or poorly regulated diabetes and/or leg artery obstructions (see also section 1.4.2). The disease often begins with a small lesion caused for example by a pinching shoe. Insensitive or limitedly sensitive to pain, the patient often does not notice the wound and the disease proceeds. If not treated in time, DFS quite often leads to gangrene (necrosis) and lower limb amputation. [34]

Figure 1.2 shows some example wounds of DFS.



Figure 1.2: Wound examples of diabetic foot diseases.

1.4.2 Peripheral Artery Occlusive Disease

Peripheral Artery Occlusive Disease (PAOD), Peripheral Artery Disease (PAD) and Peripheral Vascular Disease (PVD) are different names for the narrowing and obstruction of arteries that supply blood to the limbs. The disease is mainly caused by atherosclerosis. In most cases legs are affected. As the major risk factor for PAOD is smoking – many

content substances of tobacco harm vessels – the disease is colloquially also referred to as *smoker's leg*. Other important factors that contribute to the development of PAOD are diabetes mellitus, hypertension (high blood pressure), a persistently high cholesterol level and hyperhomocysteinemia (abnormally high level of homocysteine in the blood). An obstructed blood flow potentially results in a painful leg, ulcers as well as wounds that do not or merely poorly heal. In the worst case there occur necroses which frequently lead to the need for amputation of the affected foot or leg. [3, 4]

Figure 1.3 shows some example wounds of PAOD.



Figure 1.3: Wound examples of peripheral artery occlusive diseases.

1.4.3 Venous Leg Ulcer

Venous leg ulcer, also referred to as *Ulcus Cruris Venosum (UCV)*, is a mostly open, often painful, chronic wound on the leg or foot. Whereas PAOD concerns arteries disturbing the supply of limb cells with oxygen and nutrient, UCV affects leg veins and thus the drainage of the lower extremity. Veins have the essential task of transporting de-oxygenated blood and metabolic waste collected from cells back to the heart. Larger veins possess valves

enabling the transport of blood in one direction and against gravity. Due to insufficient valves, caused for example by an inflammation or a venous thrombosis, the deep veins of a leg and connected superficial venous capillaries are exposed to unduly increased pressure. In further consequence capillaries are damaged and leak, which lead to the development of ulcers in surrounding tissues or to a poorly healing wound in case of traumatic events. Venous leg ulcers usually arise on the inside of the lower shank above the ankle. [17]

Figure 1.4 provides some examples of venous leg ulcers.



Figure 1.4: Wound examples of venous leg ulcers.

1.5 Outline

The remainder of the thesis is structured as follows. In section 2 standard image based wound management tools in the health care sector are outlined and an overview of current image completion techniques is given. In section 3 requirements placed on an image acquisition system for wound imaging at a hospital in general and the deployed imaging set up in particular are explained. Furthermore, used reference markers for the sake of

scale and colour normalisation as well as the newly developed marker detection procedure are presented in that section.

As the developed wound course simulation tool presents the main contribution of this thesis, the next chapters are dedicated to its primary function blocks: wound analysis (section 4), wound synthesis (section 5) and wound development (section 6). Accordingly, section 4 describes a wound and its structure in general. Furthermore, several image based features are tested for its suitability in wound analysis and, as a result, proper features for wound development stage and tissue classification are presented. Section 5 describes the generation of a synthetic wound as well as deployed visual effects and section 6 explains the actual wound evolution algorithm and presents some simulation results for considered diseases. Finally, section 7 draws a conclusion, discusses the limitations of the current approach and suggests areas of future research.

Chapter 2

Related Work

The documentation of therapeutic and nursing measures on patients is regulated by law in many countries. This also applies to the therapy and treatment of chronic wounds. In addition to legislative aspects, a good documentation of the performed therapeutic measures and the wound development is also essential in the sense of a continuous improvement process in the context of quality management of a clinic or health care facility. The documentation of a chronic wound and its development in the course of therapy is often difficult or inadequately possible in written form. Therefore, it is obvious to augment the written description of a chronic wound by appropriate image data. Hence, there exists a wealth of work and even products employing image processing in the wound care context. Many of them focus on the documentation and assessment of wounds during the period of treatment.

Akestes [2] and Jalomed [25], for example, have developed computerised wound management solutions that, in addition to patient data management, wound anamnesis, therapy suggestion and other features, also enables image-based documentation of a wound including the possibility of a rudimentary measurement of the wound size by interactively drawing a polygon around the wound area(s). Scale calibration is accomplished by manually defining a known distance on a ruler placed in the image. Wild et al. [60] took it one step further and developed a wound documentation tool deploying a wound tissue classification based on colour information. Focussing on e-learning, [48, 49] describe a tool for browsing through a wound image database and allowing inter alia interactive analysis of wounds, their segmentation and classification. In order to circumvent limitations of standard 2D wound imaging, 3D-based approaches allowing wound depth measurement were investigated and related products were brought to the market in the last few years,

like the new mobile product of EKARE [15]. Considering the invisible spectrum of light, there exist wound imaging tools, like that of Kent Imaging [27] presented in [31], deploying Near Infrared Spectroscopy (NIRS) for measuring the oxygenation of the wound region which allows conclusions about the expected wound healing.

To the best of the author's knowledge, there exists no software tool predicting and simulating possible wound developments in form of a video sequence by machine vision methods together with diagnostic data. So the approach of automatic feature extraction and associated training algorithms for image based wound syntheses and simulation of possible wound development is completely new, apart from the already mentioned pre-publication of parts of this work with focus on the newly introduced image completion technique optimised for the simulation of wound development [52].

As also stated in [52], learning methods for the estimation of appearances were only considered at single moments, that is, in a statistical context [6, 28, 33, 56]. Existing simulations, for example virtual assisting systems in wound surgery [53], refer to optimum primary care only, not handling defects in wound healing with long-term treatments.

For emphasizing the importance of an uniform data acquisition the work of Lange [32] has to be mentioned. In his doctoral thesis, he analyses the advantages and pitfalls of digital photography in the field of wound management and the minimal requirements to ensure stable image-acquisition. Furthermore, Lange explains and evaluates parameters affecting the depth of field as well as the influence of the illumination onto the image quality. Also differences in compression due to the use of different cameras are examined and the suitability of different rulers for wound size measurements is investigated. These evaluations strengthen the need for using a fixed camera setup as well as a target for radiometric and geometric calibration for a reliable training data collection.

The focus of this literature review was on existing work in the context of image based wound analysis, synthesis and simulation of wound evolution. More relevant literature on various computer-vision based methods used in this thesis is referenced in the respective sections.

Chapter 3

Image Acquisition and Marker Detection

Contents

3.1	Image Acquisition System	11
3.2	Marker Detection	16

3.1 Image Acquisition System

3.1.1 Requirements

Colour image based assessment of wounds places several demands on an image acquisition system.

One aspect refers to scale. In order to determine the real size of a wound in a post processing step, it is necessary to place a proper reference object of known dimension in the scene. For the reference object a standard paper strip in our case comprising a 20 x 20 mm square can be used. The paper strip is referred to as marker within the rest of this thesis. To keep measurement errors small, it is crucial to place the marker close to the wound and in the plane of the wound as good as possible. In addition to this, the medical staff taking wound images has to be instructed to maintain a more or less constant distance of an arm's length between camera and marker located near to the wound. In this way the scale variation in the set of acquired images can be kept small.

Another aspect concerns illumination. As wound images have to be comparable with respect to brightness and colour, a proper source of light is needed during image acqui-

sition. A good light source minimizes the influence of varying ambient light and ensures uniform and reproducible illumination of the object of interest. As many chronic wounds are wet, there often occur distracting reflections in the image. In order to reduce such effects, polarising filters should be deployed.

Last but not least, the image acquisition system should be applicable in a hospital environment. Therefore, all involved equipment surfaces require to be resistant to disinfectant agents and the system should be robust, flexible and easily operable by physicians. In order to gain reproducible results, settings should be fixed and protected against accidental misalignment. It turned out to be helpful to provide a short guideline containing instructions and relevant camera settings to the image acquisition staff at the hospital.

3.1.2 Camera Setup

Following setup has been used for wound image acquisition at the hospital: A digital reflex camera, Sony α 33 (SLT-A33), presented the centrepiece of the system. It was equipped with a standard 18-55 mm camera lens. The camera offers a good price-performance ratio and meets the requirements described in the previous Section 3.1.1. In order to achieve comparable results, the camera was operated using fixed settings according to Table 3.1. Detailed information regarding camera settings and the meaning of related acronyms can be taken from the user manual [54].

Parameter	Setting
Modus	S
ISO	1600
Shutter Speed	1/100
Image Quality	Fine
Drive Mode	Single Shot
Metering Mode	Spot
Flash Mode	WL
Autofocus	AF-A
AF Area	Wide
Flash Compensation	± 0
White Balance	5600
DRO	Auto
Creative Style	Std
Flash control	Pre-Flash TTL (flash must be open)

Table 3.1: Camera settings used during image acquisition in hospital.

Based on the good experience gained in earlier projects, a camera mounted ring flash light was deployed for uniform scene illumination. The ring flash light (Metz mecablitz 15

MS-1) was used together with a bounce diffuser included in the delivery. The flash was controlled wirelessly using the camera's built-in flash operated in Through the Lens (TTL) pre-flash mode, whereas the camera's internal flash acted as master and the ring flash as slave. To suppress the influence of the master flash during light metering, it was occluded by an infra-red pass filter bracket. The remaining infra-red part of the master flash was sufficient for controlling the slave. In order to reduce reflections, two opposing circular polarising filters (Hoya Pro1 Digital 58 mm) were used — one mounted onto the ring flash, the other placed in front of the camera lens. In some cases slight reflections remained which presented no problem. They were even helpful in subsequent wound analysis steps when extracting discriminant features, which will be discussed in more detail in Section 4.4. Figure 3.1 depicts the deployed equipment at the hospital.



(a) Sony α 33 (SLT-A33), digital reflex camera; © www.davidemuci.it



(b) Metz mecablitz 15 MS-1, ring flash light; © Metz



(c) Metz infra-red pass filter clip; © digitfoto.de



(d) Hoya Pro1 Digital 58 mm, circular polarising filter; © amazon.de

Figure 3.1: Equipment deployed for wound imaging at the hospital.

To further improve the comparability of brightness and colour impression of different pictures, an additional colour normalisation step was conducted. Reference colour patterns necessary for normalisation were brought into the scene by augmenting the marker with colour fields. Note that the colour marker design and the development of colour normalisation procedure itself were not part of this thesis, but was developed based on the work of [16] and [58]. During image acquisition the camera was held by the doctor, which was more convenient than using a bulky tripod. Figure 3.2 shows the image acquisition setup.

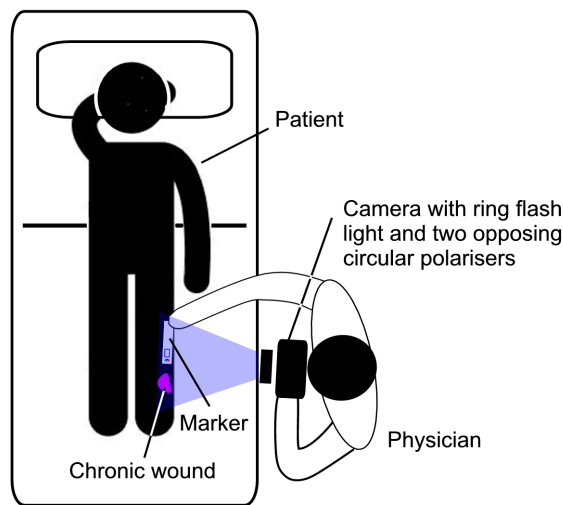


Figure 3.2: During image acquisition the patient lies on an examination table. The physician places the marker near the wound using one hand and outstretched arm. The other hand operates the camera.

3.1.3 Description of Different Marker Types


Within the scope of project SimuWound, a marker comprising a reference square and additional reference colour fields was developed (as already mentioned in Section 3.1.1). As the first marker version was improved in several steps, four different types have been deployed during image acquisition in the hospital. These versions are referred to as marker type 1 to 4 in the context of marker detection. Types 1, 2, and 3 were printed onto photographic paper in small quantity using a standard laser printer. For reasons of hygiene, a fresh marker was used each time a patient was examined by the doctor which led to a high demand for markers. So final marker type 4 was professionally produced in a large quantity in a printing plant, with the positive secondary effect of smaller inter-sample

colour differences and better colour fastness over a long period of time. Nevertheless all markers were stored protected from light.

To increase the amount of images for analyses, older wound photos showing conventional markers (four different types), were considered by the author, too. Since these markers (collectively referred to as marker type 0) merely offer a reference square and provide no colour information, they were used for investigating wound evolution and metric aspects only. Figure 3.3 shows the deployed marker types.



(a) Marker type 0 © Lohmann & Rauscher



(b) Marker type 0 © 3M



(c) Marker type 1



(d) Marker type 2



(e) Marker type 3



(f) Marker type 4

Figure 3.3: Overview of deployed marker types. a, b: Two type 0 examples - standard markers merely offering a reference square and a measuring scale used in daily hospital routine; c-f: Marker type 1-4 exhibiting additional colour fields for the sake of colour normalisation.

3.2 Marker Detection

3.2.1 Detection Approach

When solving computer vision problems based on colour images, it is often advisable to investigate different channels of several colour spaces. Selecting a proper channel (or a combination of channels) may help to reduce the complexity of a task. For example, when detecting specific patterns or objects, some channels provide better contrast than others. Following this approach, several colour spaces including Red-Green-Blue colour model (RGB), Hue-Saturation-Value colour model (HSV), L^*a^*b as well as CIE 1931 XYZ color space (XYZ) were analysed regarding its suitability for marker detection. The R plane of RGB colour space showed great promise for the detection task, since it provided the best contrast values for all marker types introduced in Section 3.1.3.

As all marker types share a blue or black reference square, the author focussed on this common element in the first step of marker detection. To solve this task, three approaches were short listed. Finding the square with the aid of Hough transformation [18] was one, deploying Maximally Stable Extremal Regions (MSER) detection [36] was another and using template matching [9] was the third approach. The first mentioned method left a lot to be desired when applied to different marker types. Due to many ambiguities in Hough space, the detection of the correct four marker square corner points proved to be cumbersome and unreliable. On the contrary, MSER detection attempts delivered promising results over a big set of test images. For a short explanation of MSER detection the interested reader is referred to section 3.2.2. In consideration of the good outcome, the work on the MSER approach has been proceeded. Hough transformation and template matching were no longer pursued.

As described in section 3.2.2, MSER+ detection delivers a set of regions that are surrounded by a dark(er) border. Applied on wound images, the white marker square framed by its continuous black or dark blue border is most likely part of that region set. In order to extract the desired marker square from the set several selection steps have to be conducted. To start with, all regions that touch the image border are discarded since marker square as well as colour fields have to be fully visible in the image. For all remaining regions, properties are calculated. These properties include *Area* (the size of a region in pixels), *FilledArea* (the size of a region when its wholes are filled), *Solidity* (value specifying the proportion of the pixels in the convex hull that are also in the region), *MajorAxisLength* (value specifying the length in pixels of the major axis of the ellipse that

has the same normalized second central moments as the region), *MinorAxisLength* (same as *MajorAxisLength* but for the minor ellipse axis) and *BoundingBox* (a vector defining the smallest rectangle containing the region).¹ Based on the calculated properties, a region has to fulfil six criteria (see inequations 3.1) to be regarded as a marker square. Note that *ImageSize* stands for the number of pixels the input image consists of. All threshold values in these criteria have been determined exhaustively in several tests.

$$\begin{aligned}
 &Area/ImageSize > 0.001 \\
 &Area/ImageSize < 0.2 \\
 &FilledArea > 1600 \\
 &(MajorAxisLength - MinorAxisLength)/MajorAxisLength > 0.2 \\
 &(FilledArea - Area)/FilledArea < 0.02 \\
 &Solidity > 0.96
 \end{aligned} \tag{3.1}$$

If no or more than one marker square candidates are found, MSER detection routine is run again using a different stability parameter. All in all, the detection procedure may take up to three loops. Note that there exists a special case: if more marker square candidates remain, the biggest one in size will be selected when all other regions are smaller than 20 % of its size. This covers the situation when the marker square together with some (smaller) colour fields remain in the detected region set. Figures 3.4(a)-(d) depict some intermediate steps of the detection procedure.

Once a proper candidate is found for the square, its four corner points are determined. This is done by applying Harris corner algorithm [18, 20] to the binary mask of the detected marker square mask, which leads to a Harris corner measure matrix, represented by a grey value image. If we now determine the location of the four highest local maxima in that image using non-maxima suppression, we will have found the exact location of the marker corners. The result can be seen in Figure 3.4(e). Exact corner points are necessary, since they define the area of the square. The ratio of this area and a known reference marker size, in turn, defines the scale of the image, which subsequently allows the calculation of the real wound size in a photo.

Given the reference geometry, the corner points also build the basis for determining the

¹The numerical computing environment MATLAB[®] [37] provides a method that calculates these and other properties out of the box. For more information, see the Image Processing Toolbox[®] documentation [55, function 'regionprops'].

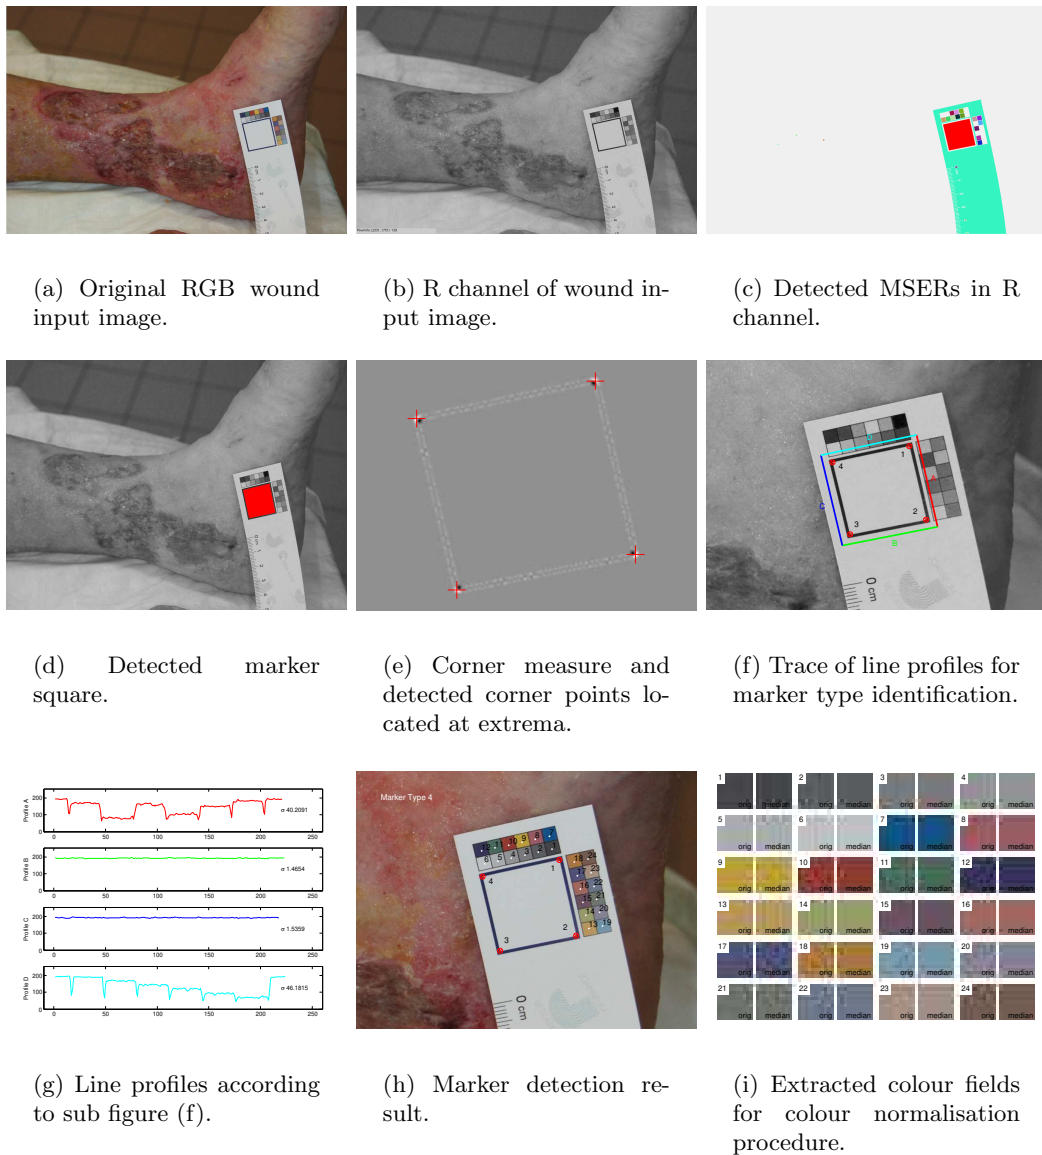


Figure 3.4: Intermediate steps of marker detection procedure.

colour field locations on the marker. As there exist different marker types offering colour fields at different locations as described in Section 3.1.3, the detector has to identify the marker type before. A simple and effective way to accomplish this task is introducing line profiles. For each line profile a standard deviation σ_i is calculated. If all four standard deviations are low the marker is of type zero, which means that no colour fields are available. If just one σ value is high, marker type 1 is existent. In case that two neighbouring line profiles exhibit a higher standard deviation than the remaining two profiles, and if

there is no gap between marker square and colour fields (determined via a short extra line profile), marker type 2 will be present. Otherwise type 3 or 4 are manifest. The distinction between type 3 and 4 is made by comparing their colour field bars because one bar is flipped. Figures 3.4(g)-(h) show the line profiles and its traces based on a marker type 4 example. Given the detected corner points and the marker type dependent reference geometry, the colour field centre positions of the marker can be obtained based on projective transformation. The 2D-homography was calculated with the aid of Direct Linear Transformation (DLT) according to [21].

The calculated colour field centre position results can be seen in Figure 3.4(h). Additionally, Figure 3.4(i) shows patches extracted from all centre positions (left column sides) as well as their median values (right column sides) which were used in subsequent colour normalisation steps. Finally, the state flow diagram in Figure 3.5 summarises the main steps of the complete marker detection procedure.

3.2.2 Maximally Stable Extremal Regions (MSER)

Originally introduced to solve wide-baseline stereo problems, MSER is not only qualified for finding reliable region based correspondences in two images, but also for detecting bright or dark bordered regions in many other applications.

Basically, MSER detection works on the topology of intensity value images. The values are best interpreted as height information. Accordingly, dark values represent lower and bright values higher terrain. After the complete flooding of the fictive grey value scene and the subsequent successive lowering of the water level (via thresholding), at some point, islands representing local grey value maxima will appear and grow. When observed from a bird's-eye perspective, the cross Section of an island at each discrete threshold level is referred to as extremal region in the context of MSER detection. With the aid of a so called component tree, extremal regions and their size are tracked during stepwise lowering the threshold value. Of course, regions will merge during this procedure until the whole image is absorbed. In the end, those regions were considered as extremal and maximally stable, that fulfil a defined stability criteria over a certain range of threshold levels. In simplified terms, those regions were picked that, more or less, kept their size during several threshold steps.

In that way, the algorithm is able to detect bright regions surrounded by a dark border or vice versa, if the input image is initially inverted. First mentioned method is referred to as MSER+, the second is defined as MSER-.

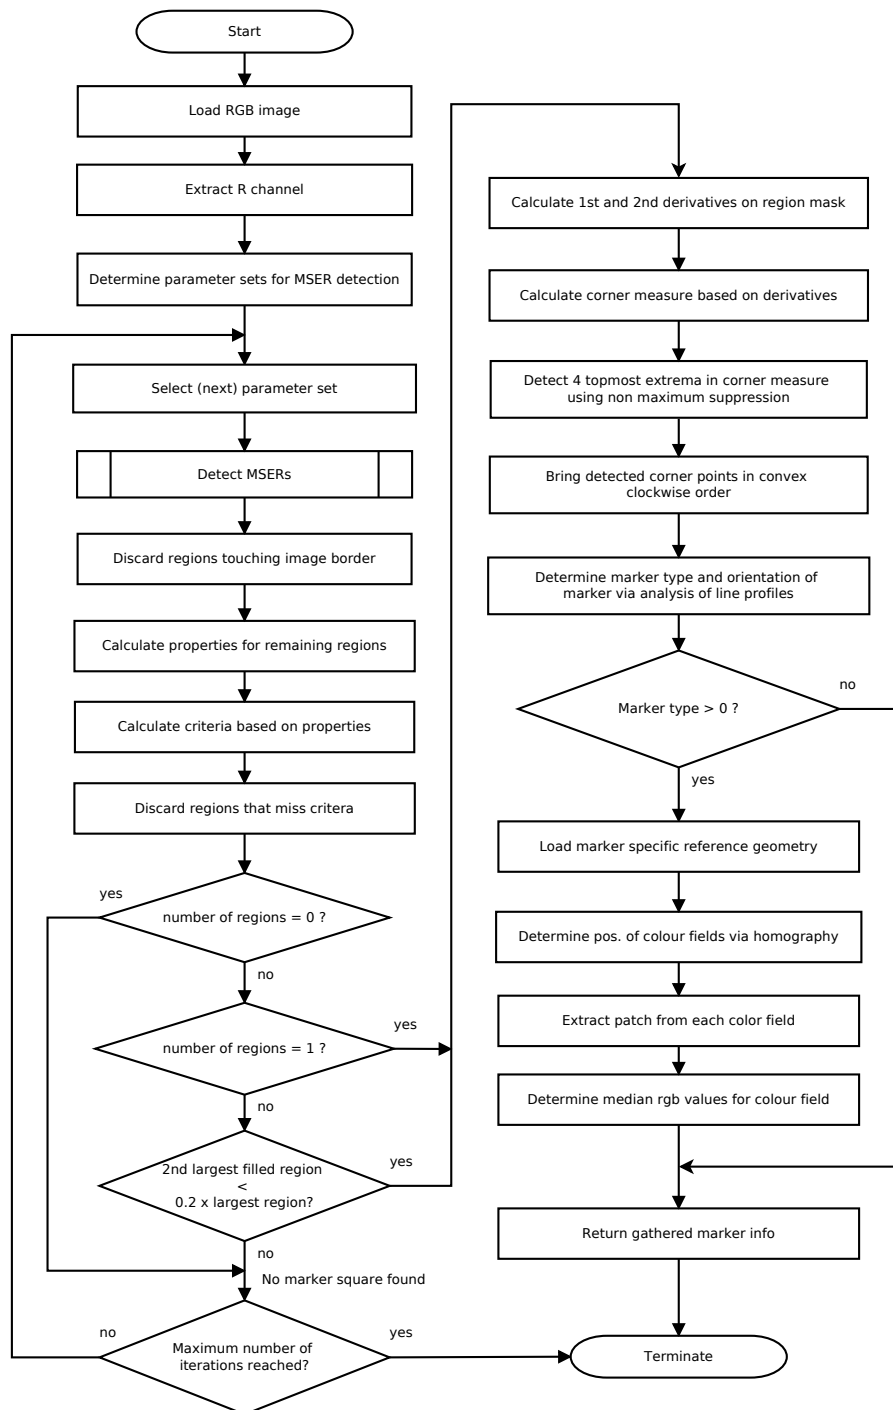


Figure 3.5: State flow diagram of marker detection procedure.

3.2.3 Performance Evaluation

The detector was designed and trained based on a small set of images. The next step is to prove that it performs on other wound images as well. Therefore, a validation experiment covering the generalisation aspect shall be set up. The test set for the experiment should comprise all available wound images: photos depicting markers with colour fields (type 1-4) captured during the project and also older images of various resolutions showing basic markers of type 0. The latter were taken with various cameras, differing from that described in section 3.1.2. Due to different image resolutions ranging from 640 x 680 to 4592 x 3056 pixels, the detector's ability to deal with scale variations can be tested inherently, too. Finally, also wound images comprising no markers (referred to as class 'none') should be part of the test set. Table 3.2 provides information about the composition of the image set deployed for performance evaluation. Note that all images are labelled by their ground truth based on six classes: marker type '0', '1', '2', '3', '4' and 'none'.

Marker type	none	0	1	2	3	4	Total
Sample images	93	136	62	100	193	105	689

Table 3.2: Image set for detector evaluation. The set consists of wound images comprising marker classes '0' to '4' as well as images showing no marker which were referred to class 'none'.

What else should be covered by the experiment? As explained in section 3.2.1, the detector has to fulfil three main tasks: detecting the marker square in the image in order to determine the proper image scale is the first, identifying the right marker type presents the second, and locating possible colour fields for the sake of colour normalisation is the third task. So, when verifying the detector's performance, these tasks should be evaluated as well. Verification of the detected marker square and marker class is trivial, it can be done by comparing the results with a ground truth. There remains the evaluation of detected colour field positions. This can be accomplished by investigating extracted colour patches in a chart. Each chart column consists of 24 colour fields. In that way wrongly detected colour field positions lead to noticeable pattern differences, which are easy to count with the naked eye. In that way, cumbersome determination of colour field centre ground truth for several hundred images are prevented.

The outcome of the evaluation can be seen in Table 3.3 and Table 3.4. The former provides information about the detection rates of each marker class, the latter shows the related confusion matrix. If all marker types were detected properly, the matrix would

Marker type	none	0	1	2	3	4	Total
Detection rate	97,8%	94,1%	98,4%	99,0%	96,9%	90,5%	95,9%

Table 3.3: Marker detection rates for each marker class (type 0 to 4) as well as for images without markers (class 'none').

		Detected Class (Marker Type)					
		none	0	1	2	3	4
Actual Class (Marker Type)	none	91	2	-	-	-	-
	0	8	128	-	-	-	-
	1	1	-	61	-	-	-
	2	1	-	-	99	-	-
	3	5	-	-	1	187	-
	4	9	-	-	1	0	95

Table 3.4: Confusion matrix indicating the detector's ability of recognising the correct marker type. Class 'none' means that there is no marker in the image.

only possess entries in the diagonal. Hence, off-diagonal entries indicate misdetections. What we can see is that the overall detection rate is 95.9 %. So, the detector works, by and large, quite well. Nevertheless there occur some detection errors in the first column of the confusion matrix, which means that in these cases a marker is present in the image but was not detected. A closer look at these particular images reveals the reasons for most of the misdetections: image blur and reflections. Normally, the deployed polarisers suppress such reflections. Unfortunately, it cannot be precisely reproduced what happened here. Since all affected images were taken more or less in the same period of time it is probable that some camera/filter settings were wrongly adjusted. In two other cases the marker square detection was disturbed by a date/time overlay in the image. Since the overlays merely occur in an older image series, it presents no problem. Figure 3.6 shows examples of images that led to misdetections.

Finally, a colour field patch chart example generated during the evaluation experiment is depicted in Figure 3.7. The chart presents the results of 50 detected markers. Altogether, 442 colour field based markers were tested in this way resulting in 9 such charts. An evaluation of the charts showed a detection rate of more than 97%.




Figure 3.6: Examples of wrongly classified marker types.




Figure 3.7: Evaluation of colour field localisation. The numbers above the chart represent the image Identifier (ID), those on the left and right indicate the colour field index.

Chapter 4

Wound Analysis

Contents

4.1	Wound Description	25
4.2	Dividing a Wound Evolution into Phases	26
4.3	Features for Wound Description	27
4.4	Image Based Descriptors for Wounds	29
4.5	Automatic Wound Classification	32
4.6	Finding Similar Patches	33
4.7	Classification of Wound Regions into Tissue Classes	38

4.1 Wound Description

A wound can be described in many ways. First of all the location of the wound can be globally specified with respect to the body using common anatomic terms. Beside the global topology the local structure of the wound can be described as well. The structure consists of a wound bed or wound base denoting the harmed skin region, a periwound region representing the transitional area to healthy skin and a wound border delimiting these two regions. The wound border constitutes the region where new epithelium develops. Epithelial tissue forms the outermost layer (epidermis) of the skin. In most cases the wound bed excretes some serous drainage called exudate, which can also be considered when describing a wound. Figure 4.1 depicts the local topology using the example of a wound caused by PAOD. Note the reddened periwound skin which is indicating an inflammation.

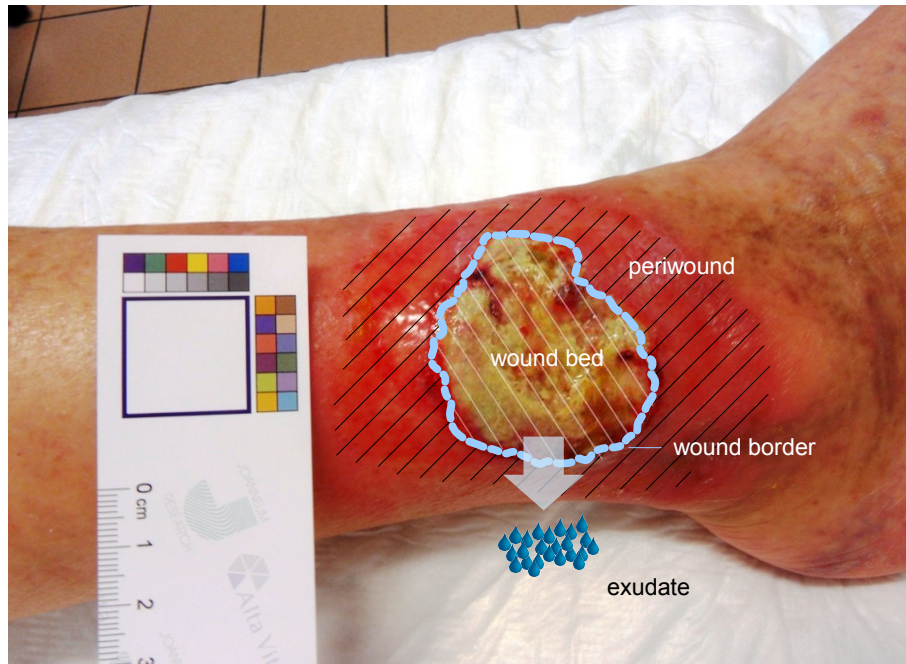


Figure 4.1: Local topology of a wound. The white hatched region represents the wound bed (wound base) surrounded by a blue dashed line - the wound border, whereas the black hatched area labels the periwound region. The redness indicates an inflammation. Blue drops indicate exudate excreted by the wound bed.

4.2 Dividing a Wound Evolution into Phases

A first attempt for the planned wound evolution simulation introduced was based on the idea that a wound passes several stages of disease during its healing or worsening process. There exist various classification schemes for wounds used in daily clinical practice. Unfortunately none of them are directly usable for automated image based processing with focus on wound simulation. For this reason the project team tried to manually classify the wound evolution process with respect to the particular disease into five stages. Phase zero represents more or less intact skin and phase four a fairly advanced stage of disease. Due to ethical considerations the team has refrained from introducing an additional phase including died off limbs exhibiting extensive necrosis ready for amputation.

So, in a first step, all pictures of a disease that were already available at the beginning of the project and which were taken from several patients, were brought into a temporal order with respect to the wound evolution on a flip chart. In a second step, observed visual classification criteria were qualitatively analysed and the images were manually classified into one of five stages together with a physician. This task was not trivial and provided

for ample discussions since there exist a great variety of different wounds resulting in a high intra class variability. Even our experts were not always in agreement. Figure 4.2 shows the final classification result based on the example of UCV.

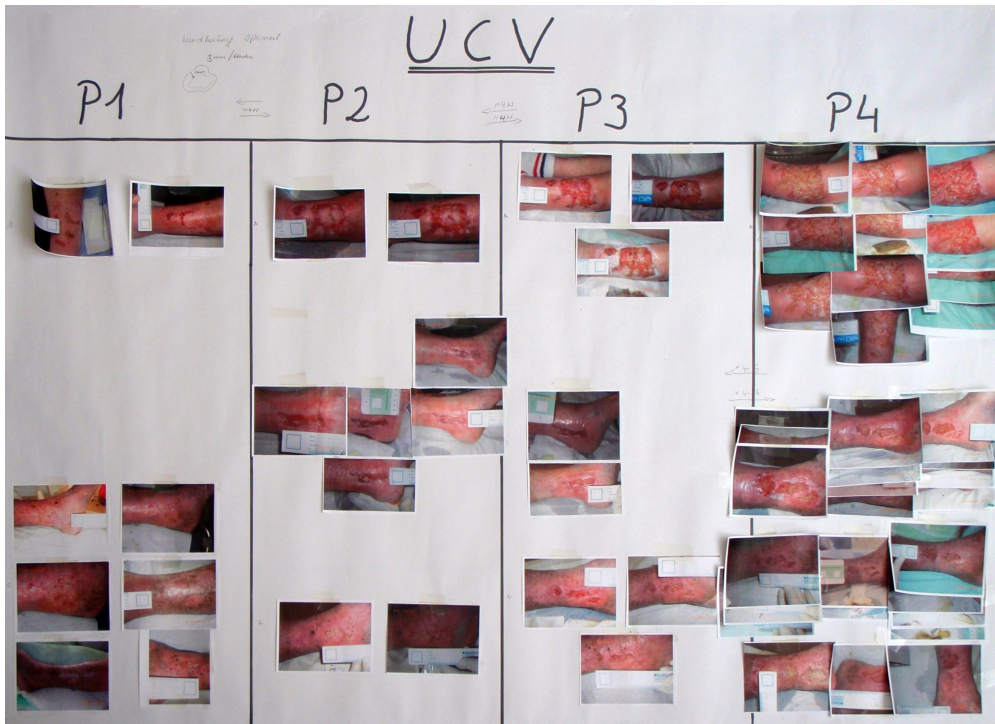


Figure 4.2: Flip chart of UCV wound images brought in a temporal order and grouped into four disease stages. Note that there is a fifth stage, healthy skin, which is not depicted here.

4.3 Features for Wound Description

After classifying all wound images into phases of wound evolution, we inspected the images once more and tried to find features for describing the wound. The most promising features from an image processing based point of view were short-listed in intense discussion with and agreement by physicians. Criteria not usable for image processing, such as the amount/nature of the exudate or smell, were disregarded. Table 4.1, 4.2 and 4.3 summarize the selected features with respect to disease and stage. For each wound evolution phase an example wound image is provided in the tables as well. Note that there exists a fifth stage (P0), healthy skin, which is not depicted here.

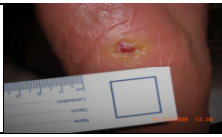
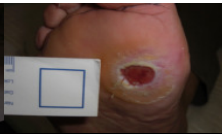


Diabetic Foot Syndrome (DFS)					
	Phase 1	Phase 2	Phase 3	Phase 4	
Example image					
Wound bed	Size/growth *	average growing	average growing	average growing	little growing; quite constant size; grows into deeper tissue layers
	Texture *	quite homogenous	granulation patterns	small cavity in wound center; granulation pattern	ragged
	Colour *	pink-red	red	red with dark regions in center of wound (small cavity)	red and yellow; extensive dark regions in the center of wound
	Shape	compact, convex	compact, convex	compact, convex	compact, convex
Wound border	small bright border	mostly distinct whitish border	mostly distinct whitish border	mostly distinct whitish border	
Periwound	normal, without findings	normal, without findings	swollen; in some cases slightly reddened	noticeably swollen; in some cases slightly reddened	

Table 4.1: Classification of DFS wounds into phases and description of visual distinctive characteristics. *-marked features were considered for automatically image based phase classification.





Peripheral Artery Occlusive Disease (PAOD)					
	Phase 1	Phase 2	Phase 3	Phase 4	
Example image					
Wound bed	Size/growth *	average growing	above average growing	little growing; quite constant size; grows into deeper tissue layers	little growing; quite constant size; grows into deeper tissue layers
	Texture *	quite homogenous	granulation patterns	subregions visible; fibrin patterns; slight granulation patterns	marbled; fibrin patterns, slight granulation patterns and necrosis pattern
	Colour *	largely pink-red; small yellow percentage (fibrin)	red (flesh-coloured); slightly covered by fibrin (yellow)	mostly covered with fibrin (yellow); slight red percentage	mostly covered with fibrin and sanies (yellow); brown/black components (necrosis); small red percentage
	Shape	compact, sometimes slightly oblong	compact, sometimes slightly oblong	compact, sometimes slightly oblong	compact, sometimes slightly oblong
Wound border	relatively sharp border; dark red if inflamed	relatively sharp border; dark red if inflamed	relatively sharp border; dark red if inflamed	relatively sharp border; dark red if inflamed	
Periwound	normal, without findings	slightly swollen	slightly swollen; in some cases slightly reddened	swollen; in some cases slightly reddened	

Table 4.2: Classification of PAOD wounds into phases and description of visual distinctive characteristics. *-marked features were considered for automatically image based phase classification.





Venous Ulcer (UCV)					
	Phase 1	Phase 2	Phase 3	Phase 4	
Example image					
Wound bed	Size/growth *	average growing	above average growing	above average growing	slower growing
	Texture *	quite homogenous	granulation patterns and fibrin pattern depending on wound	granulation patterns and fibrin pattern depending on wound	granulation patterns and fibrin pattern depending on wound
	Colour *	largely pink-red; small yellow percentage (fibrin)	red (flesh-coloured); partly covered by fibrin (yellow)	mostly covered with fibrin (yellow); small red percentage	mostly covered with fibrin (yellow); small red percentage
	Shape	one or more distinct wounds of various shape	one or more distinct wounds of various shape; wounds slightly merge	sub wounds slightly merge; epithelial islands surrounded by wound tissue	sub wounds grow together; structure covering large body part areas
Wound border	small bright border	mostly distinct whitish border	mostly distinct whitish border	mostly distinct whitish border	
Periwound	normal, without findings	normal, without findings	swollen; in some cases slightly reddened	noticeably swollen; in some cases slightly reddened	

Table 4.3: Classification of UCV wounds into phases and description of visual distinctive characteristics. *-marked features were considered for automatically image based phase classification.

4.4 Image Based Descriptors for Wounds

In a next step, it was the author's task to find suitable wound descriptors for the short-listed wound features. To recapitulate, these have been *wound size*, *texture* and *colour*.

4.4.1 Wound Size

Finding a descriptor for the wound size is a straight forward issue. Since the wound border is already available in the form of an interactively drawn polygon by experts, its enclosed area can be easily determined. As the calculated area is in pixels, it has to be converted to square millimetres by means of the detected reference marker square of known size.

4.4.2 Texture

There exist various methods for describing texture. The following subsections short-lists some promising ones for our application: Local Binary Patterns (LBP), furthermore a local variance measure often used together with LBP, a fast and easy to calculate Binary Robust Independent Elementary Features (BRIEF) descriptor. The colour aspect is covered in Section 4.4.3.

4.4.2.1 Local Binary Patterns (LBP)

The LBP descriptor is used to describe texture and has become very popular in recent years. It was originally introduced in [41]. Applied to an image, LBP provide a statistic of existing micro patterns in a texture. It is an excellent measure of the spatial structure of local image texture. The main advantage of this descriptor is its grey-scale invariance and hence also the invariance against (global) illumination variations. Figure 4.3 explains the functional principle of the basic LBP descriptor.




Figure 4.3: Principle of basic local binary patterns according to [46]. a) Local pixel neighbourhood. b) Thresholded neighbourhood. c) Binomial weights. d) Resulting LBP number of current texture unit.

According to [46], the neighbourhood of a pixel (see Figure 4.3a) is thresholded using the value of the centre pixel. The values of the pixels in the thresholded neighbourhood (4.3b) are multiplied by their corresponding binomial weights (4.3c). The resulting products are depicted in Figure 4.3d. Finally, the values of the eight pixels are summed to obtain the LBP number (134) of this texture unit. The statistical distribution of all LBP numbers of a region of interest characterises its texture.

Various extensions have been introduced for the basic LBP. [46] for example presents a rotation invariant version and [42] addresses additional multi resolution. In [43] the authors found a way to reduce the big number of different patterns to a few new significant ones, called "uniform" LBP which work well for certain applications. Since then, LBP have been used in a variety of applications ranging from face recognition [1] to paper currency recognition [19]. The survey of Pietikäinen and Zhao [47] gives a good overview of developments regarding LBP over the past twenty years.

Because of its versatility, the author decided to give the LBP approach a try to describe wound texture. The tested variants were uniform multi-resolution rotation invariant patterns referred to as $LBP_{P,R}^{riu2}$ [43], which seemed promising for our application. As it turned out, after completion of our works on SimuWound, also other authors [39] have successfully deployed LBP to describe wounds.

4.4.2.2 Local Variance Measure

The simple binary greater-or-equal grey value comparison of each pixel with its neighbours is the reason for the grey-scale invariance of LBP. However, this advantage inherently entails the disadvantage that the descriptor is insensitive to local contrast. For this reason, LBP are often used in conjunction with a local contrast measure, which is referred to as VAR in the literature (see for example [46]), and has thus established itself as a complement descriptor to LBP. The descriptor merely computes the variance over all neighbouring pixels. VAR is invariant to grey-scale shifts, but not to any monotonic transformation like LBP. A good way to combine LBP and VAR is to build a joint histogram. Parallel to the development of new LBP types described above, VAR has been further developed as well, so there exist several types in the meantime. Within the scope of this thesis the uniform rotation invariant $VAR_{P,R}$ descriptor was used for describing wound patches as well as its joint distribution together with $LBP_{P,R}^{riu2}$ referred to as $LBP_{P,R}^{riu2}/VAR_{P,R}$ (see [43]). The multi-resolution approach has been discarded because of its costly calculation.

4.4.2.3 Binary Robust Independent Elementary Features (BRIEF)

Another interesting descriptor is BRIEF [11]. In the true sense, a general-purpose feature point descriptor rather than a typical texture descriptor, BRIEF could still be of interest for describing macro texture structures in wounds, meaningfully supplementing the micro-structure-based LBP. The BRIEF descriptor uses bit strings of arbitrary length for the description of the image features, which according to its inventors are very discriminative, even with a short length. The working mechanism behind BRIEF is quite simple. Given a set of point pairs of arbitrary size and arbitrary distribution in the euclidean space, a binary intensity comparison is performed on an image for each point pair of the set. The results of the comparison tests are stored in a binary vector – the BRIEF descriptor. To speed up matching, the Hamming distance instead of the common computational expensive L2 norm can be used to compare the descriptor similarity. As a result, BRIEF is very fast both to build and to match. Per se, BRIEF is not scale invariant and rotation invariant although it tolerates small amounts of rotation. Scale invariance is no issue in our case of application since a scale normalisation is used in a preprocessing step. In [10] Calonder et al. have found a workaround to deal with rotation for database matches. A variation of this method, deploying four 90° rotated versions of the descriptor, is used within this thesis to classify wound patches and match them against a database.

4.4.3 Colour

To describe the colour of a wound and its different tissues, the author follows a common and simple approach and uses colour histograms [18]. Their main disadvantage, the susceptibility to image-capturing-related colour differences, can be minimized by a colour normalisation step we use in SimuWound and which was already mentioned in Chapter 3.1. For our case of application histograms in RGB and HSV space were tested.

4.5 Automatic Wound Classification

A wound undergoes a similar process during its healing or deterioration with respect to the underlying diagnosis. This process was divided into different phases for each diagnosis by the project partners as described in Section 4.2. In order to determine the entry point into the later-to-be-calculated wound course simulation, it is necessary to classify the wound image into such a phase in advance.

4.5.1 Cluster Analysis Based on Wound Descriptors

Based on the commonly agreed characteristics for describing a chronic wound (see Section 4.2) and the resulting image-based descriptors elaborated in Chapter 4.4, a team colleague carried out a cluster analysis [44]. The analysis should clarify whether the descriptors allow direct mapping of a wound image to a defined phase. Due to the fact that the most images were available for DFS at the time, the cluster analysis was performed on these data. Unfortunately, the analysis revealed that the amount of available data (85 wound images at that time, mostly data from single phases and not from complete wound courses) did not result in clustering in any of the feature spaces that would be suitable for classification. As an alternative to clustering, the most reliable results would probably be provided by a decision tree [8]. Based on these findings and the fact that not much more data would be available from the other diagnoses as well as the expert's know-how, the team decided to develop a decision tree approach for wound classification.

4.5.2 Decision Tree Based on Wound Descriptors

In order to decrease the likelihood of over-fitting on the existing data, good performing decision trees were sought, that are as simple as possible and possess little ramification. The implemented decision trees fulfil this requirement and work by and large well. In most tested cases (> 81%) the correct class was detected. In the final version of the wound

simulation application, the detected wound progression phase is displayed as a suggestion and can be overruled by the physician if necessary. This ensures that a possibly incorrect classification does not lead to an undesirable simulation result. Figure 4.4 shows the decision tree pseudo code for PAOD, DFS and UCV. The used wound descriptors for developed decision trees have been the total wound size, the proportion of fibrin and necrotic tissue and, in case of DFS the existence of a cavity.

```

Case PAOD
  If nNecPer > 0.03 AND (nNecPer + nFibPer) > 0.5
    nPhase = 4;
  Otherwise If nFibPer > 0.3
    nPhase = 3;
  Otherwise If nTotalPix > 10000
    nPhase = 2;
  Otherwise
    nPhase = 1;
End

Case UCV
  If nTotalPix <= 10000
    nPhase = 1;
  Otherwise If nTotalPix <= 50000
    nPhase = 2;
  Otherwise If nTotalPix <= 100000
    nPhase = 3;
  Otherwise If
    nPhase = 4;
End

Case DFS
  If Sulcus exists AND nTotalPix > 5500
    nPhase = 4;
  Otherwise If nTotalPix > 5500
    nPhase = 3;
  Otherwise If nTotalPix > 1450
    nPhase = 2;
  Otherwise
    nPhase = 1;
End

Remarks:
nTotalPix ... Total wound size in pixel
              (75x75 pixel correspond to 10x10 mm)
nFibPer     ... Relative proportion of fibrin in wound
nNecPer     ... Relative proportion of nectotic tissue
              in wound

```

Figure 4.4: Simple decision tree for classifying a wound into a phase of disease. Each diagnosis has its own tree.

4.6 Finding Similar Patches

For the wound synthesis planned in SimuWound (see Chapter 5) it was necessary to artificially create missing wound tissue or to replace existing gaps by means of suitable methods. The finally implemented method here is patch-based and requires, if there are no proper patches in the wound, suitable substitute patches. For this purpose, a database with wound patches from a variety of wound images was set up – strictly speaking, three databases, one for each disease.

4.6.1 Wound Tissue Patch Database

The patches used for the databases were taken from strictly inside a wound. There has been no overlap between patches. Furthermore, each patch database entry was provided with meta information such as the original image a patch stems from, the assigned tissue class and its scale factor derived from the marker square. In the next step, for each patch

various descriptors, introduced in Section 4.4, were calculated and added to the database. To summarize it these have been:

- Colour histograms based on RGB and HSV containing 16 bins for each channel [18]
- Uniform rotation invariant $LBP_{8,1}^{riu2}$ and $LBP_{16,2}^{riu2}$ [43]
- Local variance measure $VAR_{8,1}$ and $VAR_{16,2}$ [43]
- Joint distribution descriptor $LBP_{8,1}^{riu2} / VAR_{8,1}$ and $LBP_{16,2}^{riu2} / VAR_{16,2}$ [43]
- BRIEF-1024 (inherently including all possible shorter lengths) based on sample point distribution GII augmented to four rotated versions for 0° , 90° , 180° and 270° [11],[10]

The size of the patches in the database plays a crucial role in mentioned image-based tissue synthesis. On the one hand, too big patches lead, as experiments have shown, to coarse and 'clumsy' results. On the other hand, too small patches, in turn, may end up in unnatural, tessellated, repeating patterns. Furthermore, a too small patch size can lead to a low significance of some descriptors. For example, LBP or colour histograms are based, as the latter already reveals in its name, on histograms and sparse histograms are generally less discriminative. The conducted experimental tests reviewed by experts showed the best performance with a patch size of 25×25 pixels.

4.6.2 Matching Reference Patches Against a Database

In order to match a patch against the database a proper distance measure is necessary. For BRIEF the fast to match Hamming distance was used as suggested in [11]. And for the other descriptors three different distance measures were tested: the sum of absolute differences, the sum of squared differences, and histogram intersection [38].

As stated in a preliminary test, the sum of absolute differences and histogram intersection measure performed roughly equally well. The sum of squared differences delivered slightly worse matches in the context of similarity. So, the author has chosen the sum of absolute differences for as a distance measure. Based on the mentioned preliminary test, three of the existing descriptors in the database were short-listed. First, the combined $LBP_{16,2}^{riu2} / VAR_{16,2}$ descriptor, which performed better than its individual parts and better than the (8,1)-version. Second, the RGB colour histogram, which outperformed the HSV version and third, BRIEF-256. The longer 512 and 1024 versions proved to be overkill. They brought no detectable improvement over the BRIEF-256 version.

In a further step the remaining descriptors were tested by matching 50 randomly chosen reference patches against the database. The best five hits for each patch and each short-listed descriptor were reviewed and consensually evaluated together in the project team. In order to accomplish this task as efficiently as possible, it was necessary to prepare the representation of the results in advance in a meaningful way. So the five matching results for each reference patch and each short-listed descriptor were represented as three rows each consisting of five blocks side by side. Let us go over that with an example of representative granulation patch matching results depicted in Figure 4.5.




Figure 4.5: Matching result using a little textured granulation patch from PAOD database. 1st row: best five matches for BRIEF-256 descriptor; 2nd row: best five matches for $LBP_{16,2}^{riu2}/VAR_{16,2}$ descriptor; 3rd row best five matches for RGB colour histogram.

The first row represents the best five BRIEF-256 matches, the second row the best five $LBP_{16,2}^{riu2}/VAR_{16,2}$ based matches and the last row the best five matches for RGB colour histogram descriptor. A match is represented by a block consisting of 3x2 fields. The left three fields belong to the reference patch and the three on the right belong to the patch found in the database. The upper two fields show the respective original RGB patch, the two fields in the middle the associated grey value version and the lower two fields show a contrast enhanced version of the grey value patch, in order to better recognize and evaluate any existing textures and structures. Finally, the displayed text on top of each block reveals the ID of each patch and some information about its origin. The displayed value below each block represents the respective matching distance, the Hamming distance for BRIEF-based matches and the sum of squared differences for the others.

As we can see, the BRIEF-256 descriptor provides matches that are different in brightness and colour, and the texture in the 3rd and 4th match is also very different. $LBP_{16,2}^{riu2}/VAR_{16,2}$ performs better. A good indication of this is the origin of the patches. Four out of five patches stem from the same image as the reference patch. Merely the 4th hit differs in colour. The RGB colour histogram descriptor also provides nice matches in terms of visual similarity. Again, two matches come from the same picture as the reference patch. All in all, $LBP_{16,2}^{riu2}/VAR_{16,2}$ and RGB colour histogram score well in this example.

Now let us examine another matching result, this time for a fibrin reference patch (see Figure 4.6). At a first glance we can see that BRIEF-256 and $LBP_{16,2}^{riu2}/VAR_{16,2}$ provide unsuitable matches because of huge colour differences. The RGB colour histogram descriptor performs better, though there are noticeable texture differences for some matches.

Finally let us investigate a matching result of a highly textured reference patch (see Figure 4.7). The patch shows wet granulation tissue of a PAOD wound. Here we can see at first glance that BRIEF-256 is overwhelmed with the task. $LBP_{16,2}^{riu2}/VAR_{16,2}$ and RGB colour histogram provide some useful hits, although $LBP_{16,2}^{riu2}/VAR_{16,2}$ fails twice in terms of colour and the RGB colour histogram mostly fails in terms of texture.

After analysing the matching results for all 50 reference patches in the team, we got a similar picture as after examining the three provided examples. First, the BRIEF descriptor may work well in other tasks, but it is not the descriptor of choice for our case of application. Second, $LBP_{16,2}^{riu2}/VAR_{16,2}$ provides good results regarding texture similarity, but fails in many cases regarding colour. And finally, the RGB colour histogram provides exactly opposite results: it works well in terms of colour similarity, but fails

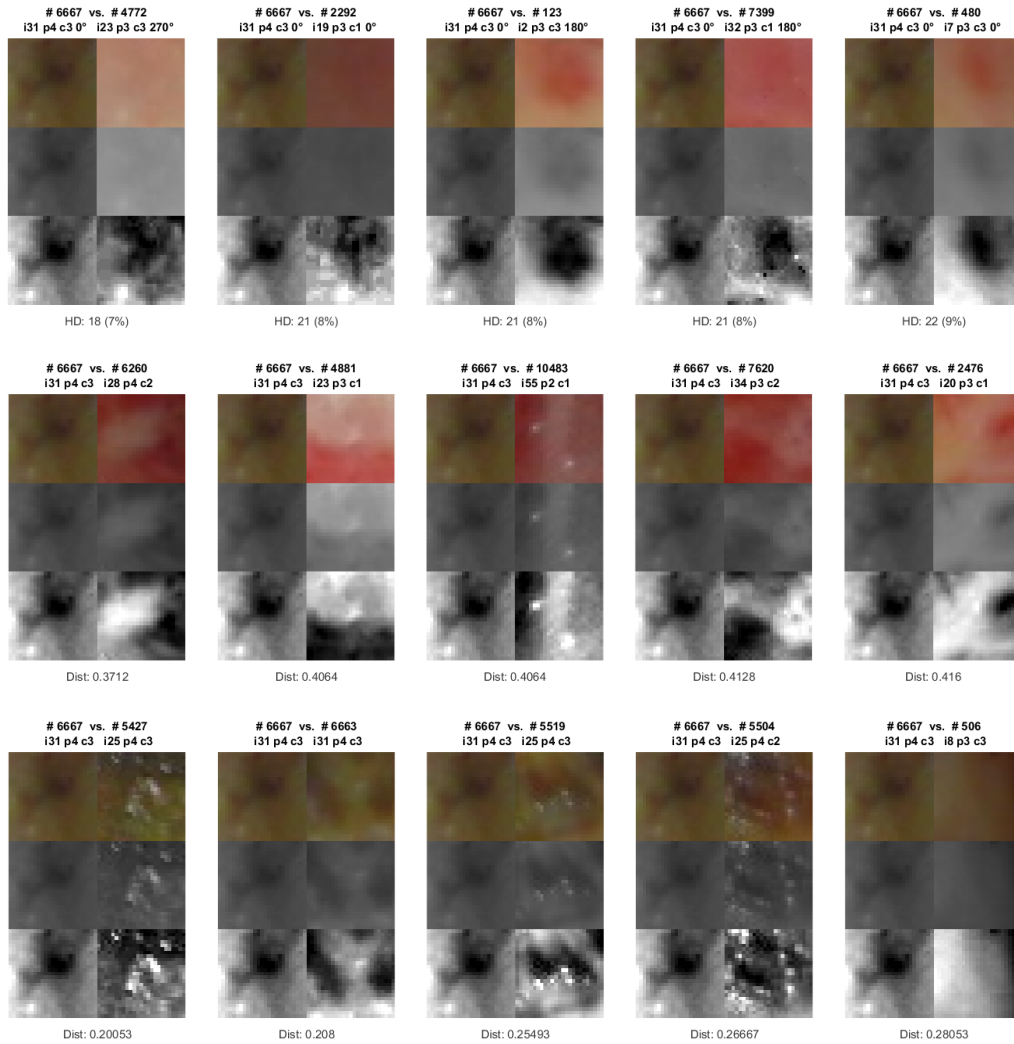


Figure 4.6: Matching result using a fibrin reference patch from PAOD database. 1st row: best five matches for BRIEF-128 descriptor; 2nd row: best five matches for $LBP^{riu2}/VAR_{16,2}$ descriptor; 3rd row best five matches for RGB colour histogram.

to match similar texture. Now the question arises whether $LBP^{riu2}/VAR_{16,2}$ and RGB colour histogram should be combined. To answer this question, the author has tested the concatenated descriptors for all 50 reference patches as well. For the result see Figure 4.8. Note, that for the sake of comparison the same example reference patches were used. We can see, that the combined $LBP^{riu2}/VAR_{16,2}/RGB$ -colour-histogram descriptor does a good job and therefore was deployed for finding similar patches in synthetic tissue completion described in Chapter 5.3.

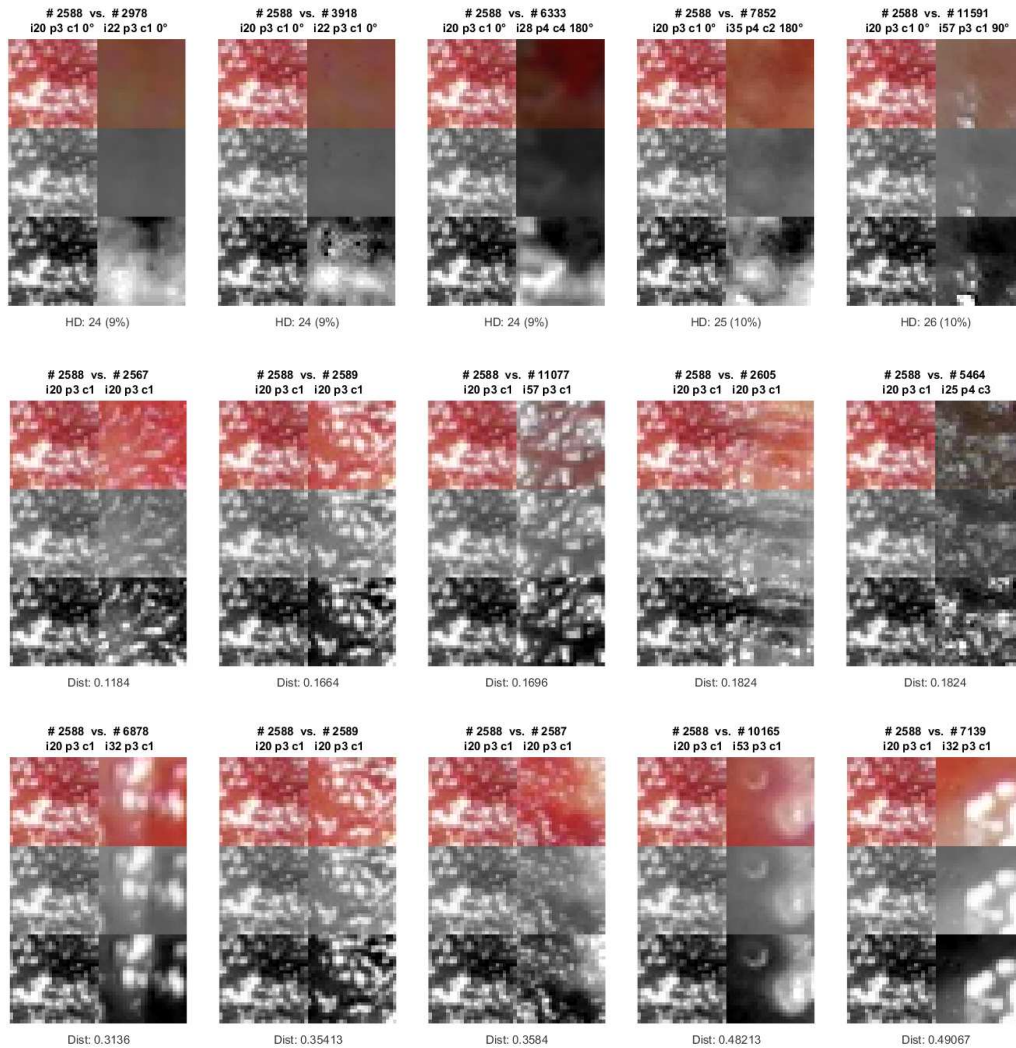


Figure 4.7: Matching result using a highly textured reference patch from PAOD database. 1^{st} row: best five matches for BRIEF-128 descriptor; 2^{nd} row: best five matches for $LBP_{riu2}/VAR_{16,2}$ descriptor; 3^{rd} row best five matches for RGB colour histogram.

4.7 Classification of Wound Regions into Tissue Classes

As already pointed out in [52] a wound physiologically consists of different tissue types depending on the stadium of disease. These main types are epithelising, granulation, fibrous and necrotic tissue. For the wound synthesis described in Chapter 5, a separation of the individual tissue types is needed, which is most easily accomplished by colour. The different wound tissues are distinguishable. According to our medical expert and to [45]




Figure 4.8: Matching result of three already tested reference patches from PAOD database, but this time a combined descriptor of RGB colour histogram and $LBP_{16,2}^{riu2}/VAR_{16,2}$ was used.

epithelising tissue possesses, for example, pink shades, granulation tissue has red shades, fibrous tissue is yellow and sometimes slightly green- or blue-tinted if special bacteria like pseudomonas are involved and necrotic tissue is normally black sometimes brown or grey but mostly dark coloured. Nevertheless, the colour based approach entails some potential sources of error. Fresh blood for instance looks like granulating tissue and dried blood has a similar dark appearance to necrotic tissue. Therefore, and to avoid other misclassification due to wound dressings and ointments, it is necessary to clean the wound before taking an image. Furthermore some wounds especially those of DFS sometimes exhibit cavities,

which indicates an advanced stage of disease, also may be confused with necrotic tissue. To circumvent this problem the cavities have to be detected separately and the affected wound subregions have to be excluded from the colour segmentation. Since conventional region growing methods have failed in the detection of cavities, the author focused on another attempt based on Statistical Region Merging (SRM).

4.7.1 Statistical Region Merging (SRM)

SRM, introduced in [40], is an image segmentation method based on region growing and merging. The underlying model of image generation implicitly makes the assumption that observed colour variations inside a regions should reasonably be smaller than between regions. According to its inventors, SRM is also able to separate regions with high variability, which is the case for many textures like grass, or in our case, speckled skin or wounds and avoids producing holey segmentations – a drawback of many other region-merging techniques. SRM is very robust against noise corruption and is able to cope with occlusions. Running the algorithm does not require much parameter tuning, the control of a single statistical complexity parameter Q makes it possible to adjust the segmentation scale, the coarseness of the segmentation and the statistical complexity in a simple manner. The higher Q the finer the segmentation and the more regions the final segmentation result will consist of, and vice versa.

4.7.2 Segmentation of Cavities Inside a Wound

After investigating different colour spaces like RGB, HSV, $L^*a^*b^*$ and others the author selected the R panel of the RGB image (see Figure 4.9(a) and (b)). The R plane provided the best compromise between high cavity contrast and low interference of other tissue types. In a next step the SRM algorithm was applied to the image Section comprising the wound (see Figure 4.9(c)). For the statistical complexity parameter Q a value of 2^9 was used. This value performed best in several tests. In a final step, areas with lower intensity values than their neighbours were identified by iterative local minimum region detection. If such a minimum region fulfils certain size ratio and standard deviation criteria it is merged to the set of detected cavity regions. Figure 4.9(d) and (e) show the intermediate detection results. The final cavity detection result is depicted in Figure 4.9(f). The potential cavity regions outside the wound are ignored. Figure 4.10 shows some more detection results for different DFS wounds. Note that the chequered pattern in the wounds of column 2 and 3 is not of physiological origin. It stems from the used dressing material.

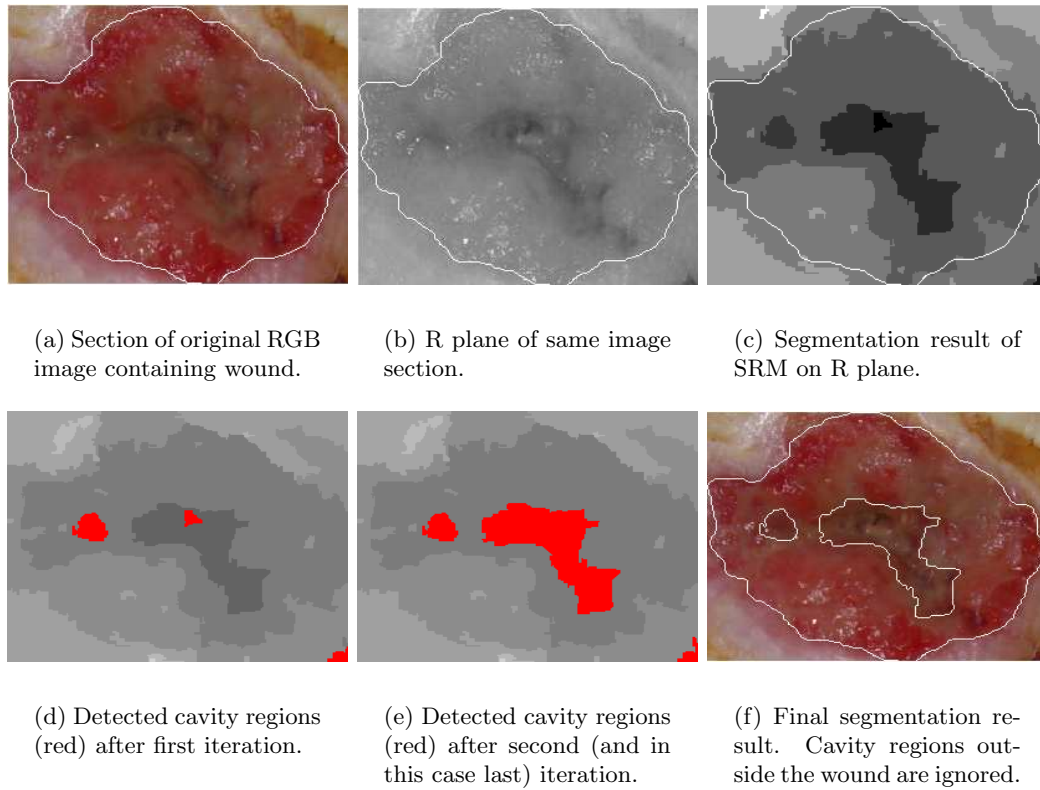


Figure 4.9: Intermediate steps of cavity detection procedure.

The cavity detection algorithm was tested on all available DFS wounds. Table 4.4 shows the results. As we can see, the detection works reliably. The sensitivity of the detection is at 82% and its specificity is at 100% for the tested images. However, the author is well aware that for an extended evaluation of the detector more than just 84 sample images would be necessary. Unfortunately this was not possible due to the limited data available.

		Ground truth	
		Wound with cavities	Wound without cavities
Detection result	Cavities detected	9	0
	No cavities detected	2	73

Table 4.4: Results of cavity detection for the tested DFS images.




Figure 4.10: Some more cavity detection results. The upper row shows the wound Section of the RGB input image. The lower row shows the respective cavity segmentation result.

4.7.3 Colour Based Segmentation

After the cavities are segmented, we can start with the actual wound tissue segmentation. Wild et al. [60] have already successfully detected different tissue types in wounds based on colour information in their developed wound documentation software. Unfortunately, they do not reveal in their work how they did it. Anyway, after renewed examination of different colour spaces, this time with a focus on the wound tissue segmentation task, the RGB space was found to be suitable.

Based on good experience in other projects and the similarity to current problem, the author decided to test a Support Vector Machine (SVM) [12] to solve the classification task. The SVM was trained on about 600 RGB samples from different tissues in 50 colour normalised wound images using a radial basis function kernel and based on leave-one-out cross evaluation according to the proposed procedure in [22]. Figure 4.11 shows the distribution of wound tissue training data in RGB space.

In a final step the trained model has been tested on 25 images not used during the training. The segmentation result was reviewed by wound experts and found to be good and valid. An example segmentation can be seen in Figure 4.12. The left image shows a colour normalised test input image and the right image the resulting segmentation.

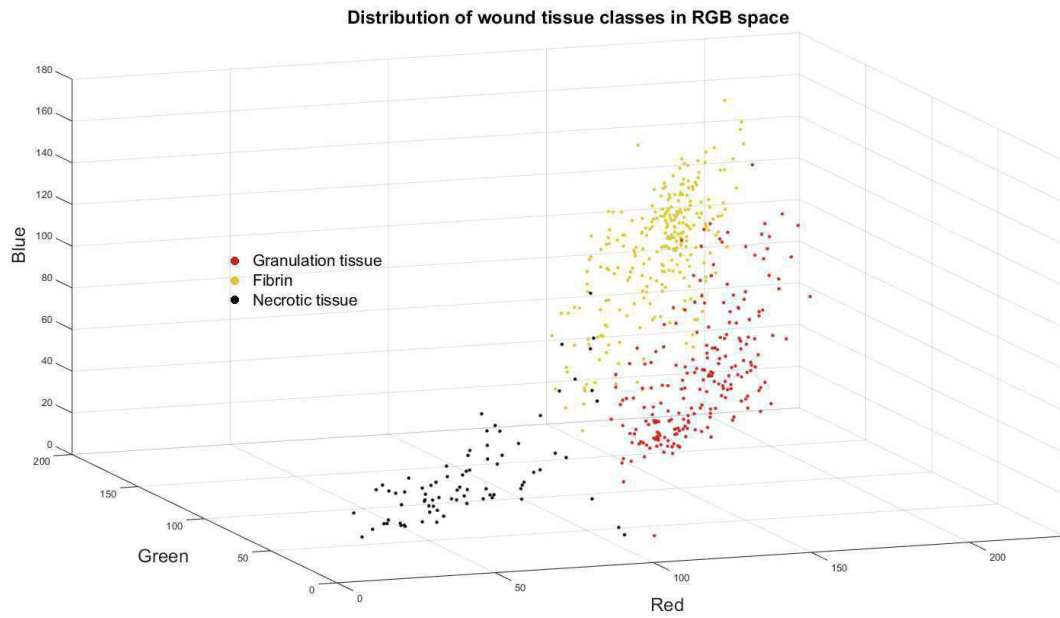



Figure 4.11: Distribution of wound tissue training data in RGB space.



(a) Section of colour normalised RGB input image containing wound.

(b) SVM based colour segmentation result.

Figure 4.12: Example result of SVM based colour segmentation in RGB space.

Chapter 5

Wound Synthesis

Contents

5.1	Skin Segmentation	45
5.2	Tissue Layer Concept	53
5.3	Generating Synthetic Wound Tissue	56
5.4	Completion of Wound Areas with Healthy Skin	56
5.5	DFS Wound Border Treatment	58
5.6	Wound Cavity Synthesis	63

5.1 Skin Segmentation

To achieve a realistic appearance of an artificially created wound in an image, first of all, it is necessary to ensure that the wound does not grow beyond the affected body part. Therefore, a skin segmentation step is essential. In order to spare the physician the tedious manual segmentation of the body part, it was necessary to develop an automatic segmentation that fits into the overall wound synthesis/simulation process.

The study of Saxen and Al-Hamadi [50] gives a good state of the art overview about colour based skin detectors and their performance. The work distinguishes between three general methods: model-based, threshold-based, and region related ones. Based on the study and the fact that colleagues have already made first good experience with it, the author's selection was the promising model-based approach from [26]. A new second approach, which falls into the region growing class, was also developed as contribution within this thesis. Both methods were evaluated for their suitability for actual application

in a straight forward manner. Details on this and a comparison of the results can be found in Table 5.1 at the end of this section.

5.1.1 Model Based Skin Detector

Jones and Rehg [26] introduced a Gaussian mixture model for skin and non-skin classes from a huge labelled pixel dataset, with the intent to detect naked people in images. Their classifier relies on colour information and operates directly on pixels.

In order to test the performance of this Gaussian mixture model and its possible suitability for our application, an appropriate framework was built. Figure 5.1 describes the skin detection framework step by step. Applied on a wound image (see Figure 5.1(a)), the detector delivers a skin probability map (see Figure 5.1(b)) and a no-skin (background) probability map (see Figure 5.1(c)), both according to Jones and Rehg. The brighter a pixel, the higher the probability for the presence of respective class. The next step is binarisation of both maps using proper threshold levels. In conducted tests, a value of $2.5 \cdot 10^{-7}$ for the skin class and a value of $2 \cdot 10^{-6}$ for the background have been found to be adequate. Figure 5.1(d) depicts the binary results of the skin map and Figure 5.1(e) that of the background map. At this point we can clear all pixels in the skin mask, that are set in the background mask, which results in an adjusted skin mask (Figure 5.1(f)). As we have additional skin knowledge from wound annotation in the form of a wound mask – a wound is always part of the skin – we merge the adjusted skin and the wound mask obtaining a new intermediate skin mask (see Figure 5.1(g)). This mask yet has many micro and macro holes. Morphological closing (see Figure 5.1(h)) and filling of (false) background regions surrounded by skin pixels close most of the holes (see Figure 5.1(i)). As we can see, there are some small skin artefacts left. We can get rid of them by selecting only the region containing the known wound area (see Figure 5.1(j)). A final smoothing step by means of Gaussian filtering and consecutive binarisation at threshold value 0.5 generates the final skin mask depicted in Figure 5.1(k). The last Figure 5.1(l) shows the skin segmentation result in the original image.

The described skin segmentation procedure has been tested in two experiments. The first using 107 original wound images and the second using the colour normalised version of the images. Although the detector gives surprisingly good results in many cases (see Figures 5.1(l) and 5.2(a)), it is sometimes over-challenged, mainly due to the applied fixed threshold when generating the skin mask from the probability map. This effect, for example, can be observed in Figure 5.2(b), where the detector responds to shadows cast

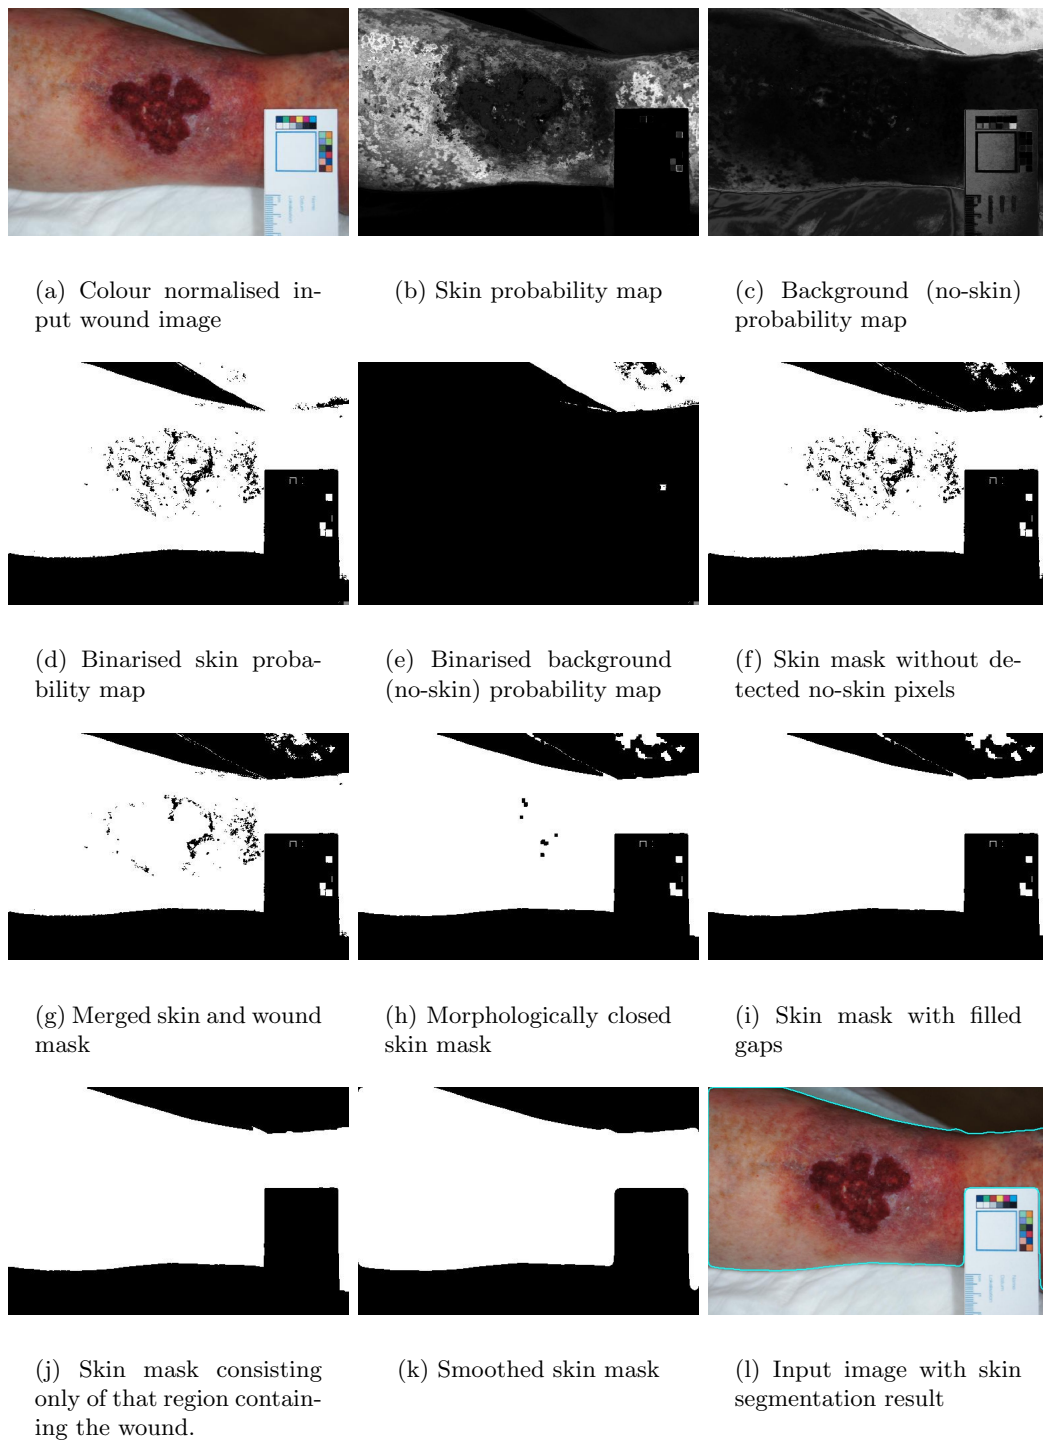


Figure 5.1: Intermediate steps of a skin detection framework deploying a Gaussian mixture colour model presented by [26].

by a foot or in Figure 5.2(c) due to the pale skin in the foot region. Nonetheless, many non-perfect segmentation results are suitable for further use as long as they are correct in the wound region. Detailed performance results can be found in Table 5.1.



(a) Good skin detection result (above) and related skin probability mask (below).

(b) Poor skin detection result (above) and the related skin probability mask (below). The detector wrongly responds to shadows cast by a foot.

(c) Borderline skin detection result (above) and related skin probability mask (below).

Figure 5.2: Some skin detection results on wound images produced by a framework deploying a Gaussian mixture model [26].

5.1.2 Region Growing Based Skin Detector

Since SRM behaved well for the segmentation of cavities inside a wound as described in Chapter 4.7.1, the method was also applied for skin segmentation. As already explained there exists only one parameter Q for controlling the coarseness of the segmentation and the statistical complexity of SRM. To find a proper value the SRM algorithm was applied to 107 colour and size normalised wound images using different Q values and the segmentation results were investigated. Figure 5.3 shows the outcome of this approach illustrated on one example wound image.

On the one hand, the statistical complexity Q should be chosen high enough so that the related segmentation is fine enough that no region includes skin and background pixels. When looking at Figure 5.3, we can see that this is the case for Q levels $\geq 2.5^5$. On the

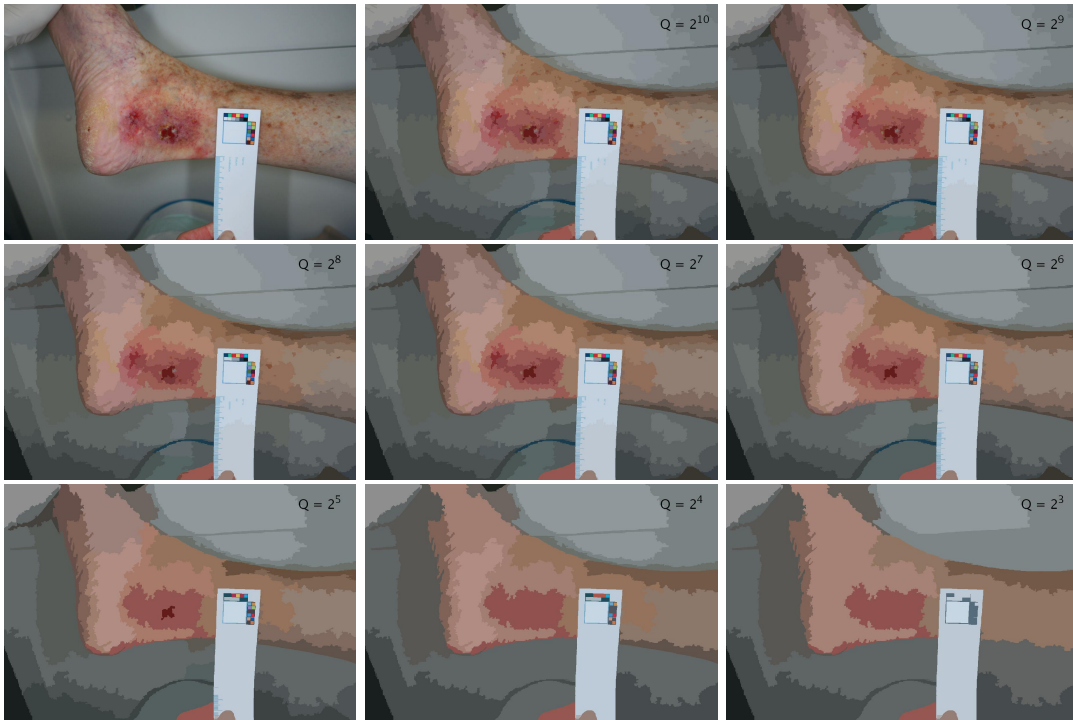


Figure 5.3: Results of SRM segmentation applied to a wound image using different statistical complexity parameters Q starting from high values to smaller ones. The upper left image shows the colour and scale normalised input image.

other hand, the segmentation should be as coarse as possible for the sake of complexity reduction. Q levels $> 2^7$ end up in too fine segmentations. After screening all SRM segmentation results the author regarded a Q level of 2^6 for a good compromise for skin segmentation in our case of application.

The next step in obtaining a final skin mask is the assignment of the correct class (foreground/skin or background/no-skin) to each segmented region and merging it according to its class affiliation. In order to automate this procedure by designing and applying a proper model we need a reliable ground truth. Therefore, 1032 SRM regions stemming from 53 training wound images were labelled interactively using a specially created GUI. Figure 5.4 shows the distribution of all labelled regions in RGB space, with red dots indicating skin and blue dots representing background regions. As we can see, skin and background class samples occupy more or less disjoint regions in the RGB space, which means that there exists a good chance of separating the two regions. The HSV space was also examined, but did not provide better clusters. Region separation in RGB space

can be done in different ways. Three methods were implemented and investigated: manually adjusted constraints (thresholds and colour ratios), a SVM [12] and the already used Gaussian mixture model approach described in Section 5.1.1, but this time applied to regions instead of pixels. Figure 5.5 shows some SRM segmentation examples using SVM as classifier.

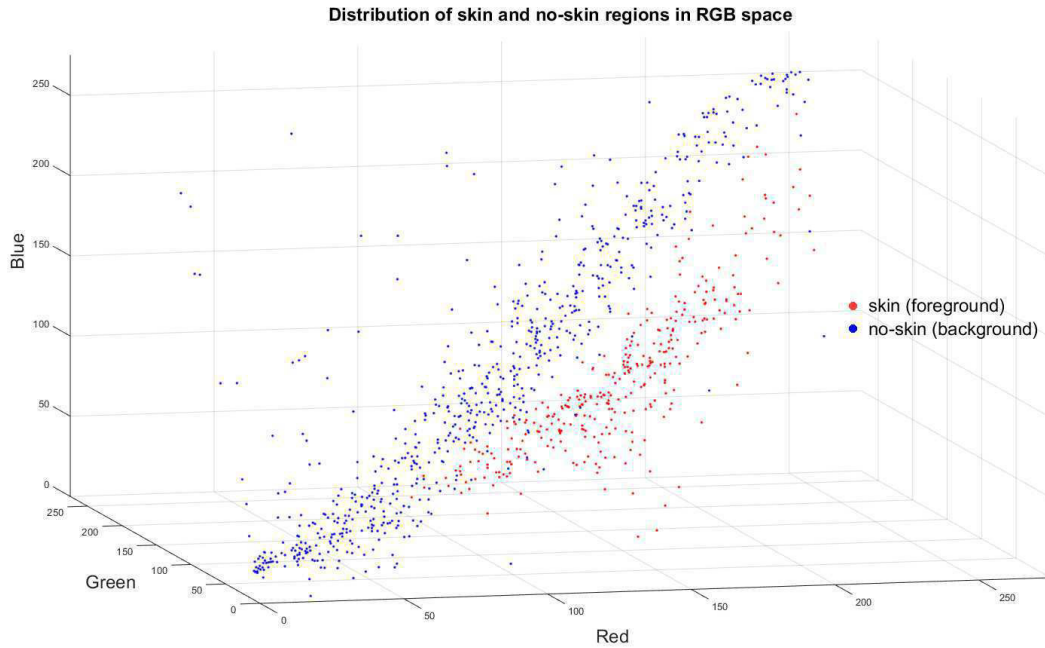


Figure 5.4: Distribution of skin and background training data in RGB space.



Figure 5.5: Exemplary skin detection results produced by SRM and using a SVM as classifier. The images have been colour normalised in a pre-processing step.

Of course, in order to make a statement about the performance, a few examples are not enough, so all three methods were tested in two experiments in the same way as described

in Section 5.1.1. For detailed performance results, see Table 5.1 in the next Section 5.1.3 which also explains the test criteria.

5.1.3 Performance Evaluation

The next step is to find a good method to evaluate the performance of the discussed detectors. One common way would be to provide all test images with a ground truth by labelling pixels of skin regions as foreground and those of non-skin regions as background and then calculate the percentage of correctly and incorrectly detected pixels for each class. Unfortunately, this method would merely be conditionally suitable for evaluating the ability for our application. Since we use the detector to find out the boundary of the body part in the image to prevent a simulated wound from growing over skin areas into the background, misclassified pixels somewhere near the edge of the image are usually less tragic than pixels close to the wound. Also, for example, misclassified isolated pixels are usually less disruptive than connected ones. Another problem lies in the definition of the mentioned ground truth itself. How should we deal with images containing skin far behind the focal plane, because, for example, the face or the arm of a patient is also photographed? This is especially true for DFS patients whose sole need to be photographed (see for example Figure 1.2). Should such regions be part of the ground truth? If so, and the detector fails to detect such fuzzy and mostly dark regions, it receives a bad sensitivity score. In the opposite case its specificity rating would be bad.

For these reasons, the author decided to review each segmentation result and to evaluate it qualitatively with respect to the suitability for our application. The assigned ratings were '*Very good*', '*Suitable*' and '*Not suitable*'. Figure 5.6 shows reference segmentation samples for each rating. Furthermore, Table 5.1 gives a performance summary of the tested detectors in absolute numbers and the two charts in Figure 5.7 show the detector performances in percentage. The upper chart refers to original wound images and the lower chart to their colour normalised versions.

Taking a look at these results, it can be seen that the region based SRM methods outperform the pixel based Gaussian Mixture Model (GMM) approach and SRM in combination with SVM classification produces the best results of all tested methods. Furthermore, it can be observed that SVM classification provides similarly good results on test and training data, which is an indication that the model is well trained and that there is no over-fitting. Finally, it can be observed that the colour normalisation preprocessing step generally increases the skin segmentation performance. Based on these findings, we



(a) Skin segmentation examples rated as 'Very Good'.

(b) Skin segmentation examples rated as 'Suitable'.

(c) Skin segmentation examples rated as 'Not suitable'.

Figure 5.6: Reference examples for skin segmentation performance rating.

Segmentation result rating		Gaussian mixture modell (GMM) (applied to pixels)		Statistical Region Merging (SRM)							
		Original Images	Colour normalised	Manually adjusted thresholds		Gaussian mixture modell (GMM) applied to regions		Support Vector Machine (SVM)			
				Original Images	Colour normalised	Original Images	Colour normalised	On test images		On test & training images	
Original Images	Colour normalised	Original Images	Colour normalised	Original Images	Colour normalised	Original Images	Colour normalised	Original Images	Colour normalised		
Very good	30	56	48	71	62	72	45	44	81	88	
Suitable	35	31	12	16	26	15	7	10	22	19	
Not suitable	42	20	47	20	19	20	2	0	4	0	

Table 5.1: Performance comparison of different skin segmentation methods.

decided in the project team to deploy the SRM-SVM method in our application.

One more note to conclusion. To further improve the segmentation results, the procedural instruction of wound image capturing has been amended so that extremities should only be photographed against a homogeneous background such as a bed sheet.




Figure 5.7: Performance comparison of different skin segmentation methods based on original images (above) and colour normalised ones (below).

5.2 Tissue Layer Concept

As we have learned in Section 4.7, a wound consists mainly of four tissue types and is surrounded by intact skin. Current wound synthesis concept also considers these tissue types.

5.2.1 Layer Stack

Each tissue type is represented by a dedicated layer and the layers are stacked one over another. The basis layer represents intact skin. The next higher layer contains granulation tissue and, strictly speaking, also epithelial tissue. Then a fibrin tissue layer is added and the upper most layer refers to necrosis. The whole stack can be seen in Figure 5.8. The procedure of generating the skin layer is described in Section 5.4 and that for the other layers in Section 5.3.




Figure 5.8: Stack of tissue layers for wound synthesis. [52]

Before we can start the wound synthesis using the tissue layers, we need to introduce tissue masks. One mask for each layer. The masks have the same size as the tissue layer and control the region of the dedicated tissue type. Mask entries near or equal zero block the layer, higher values up to 1 make the layer visible. The generated wound is essentially a linear combination of the layers and the mask entries represent the coefficients. In addition to the tissue masks, there exists a mask defining the region of the marker and one for the body part region which is important for limiting the extent of a growing wound in the image. In the case of DFS diagnosis, two more masks are needed. The first is associated to the wound border region, which differs from PAOD and UCV by a more or less wide bright, usually whitish/yellowish area and therefore has to be treated separately. The second mask is dedicated to possible cavities in the wound that as well must be treated separately.

5.2.2 Tissue Layer Masks

Let us take a closer look at the creation of the mentioned masks. The skin mask directly relies on the results provided by the skin segmentation procedure already introduced in Section 5.1.2. The three masks for granulation/epithelial tissue, fibrin and necrosis tissue are generated using the colour based tissue segmentation procedure described in Section 4.7.3. Before being used, the respective mask is subjected to morphologically operations and gauss filtering to make it more compact and smooth. The next mask to be explained is the marker mask. It is necessary to prevent a wound during its deterioration process from growing into the marker. The mask does not have to be calculated extra, it is a

by-product of the MSER based marker detection described in Section 3.2. An example of such a marker mask is depicted in Figure 3.4(c). It is represented by the cyan region. Of course, the holes have to be filled before its usage. Let us move on to the next mask, the wound border region mask used for the synthesis of DFS wounds. Its generation is described in detail in Section 5.5.2. Remains the cavity mask which is used to remove a cavity during healing simulation. The mask is gathered using the the cavity detection procedure explained in Section 4.7.2.

5.2.3 Creating a Synthetic Wound

Now that all layers and masks have been explained, we can turn to the actual wound synthesis according to [52]. Figure 5.9 shows the procedure based on PAOD step by step. To start with, for each tissue type a suitable pre-calculated texture was taken from a database. Before a texture is used, it is fused with the respective tissue already present in the wound (granulation, fibrin or necrosis tissue) by means of blending. Since in most cases the impression of brightness of the synthetic texture does not match that of the existing tissue, a brightness adjustment in the range of -20% to +15% is conducted. In a next step the different layers were merged one by one, beginning with the granulation layer (see Figure 5.9(a)), then fibrin layer (see Figure 5.9(b)) and finally the necrosis layer (see Figure 5.9(c)). The fusion was done by means of blending using the specific tissue related masks described in 5.2.2. After adding a synthetic wound border by darkening the wound edges (see Figure 5.9(d)) an additional shading filter mask was applied to attenuate wound regions close to the contour of the body. The shading filter mask was generated by morphological erosion of the skin mask followed by a Gaussian filtering. This causes a realistic 3D-like effect of the wound and prevents the wound from protruding the affected body part (see Figure 5.9(e)). Finally, the generated wound tissue is merged with the skin layer (see Figure 5.9(f)). The original wound (before colour normalisation) can be seen in Figure 5.9(g) and the final result of the wound synthesis at an advanced stage of disease is presented in Figure 5.9(h). In addition to the described synthesis there were more effects implemented to increase the realistic visual impression of the created wound. These are described in the context of wound development simulation in Section 6.2.




Figure 5.9: Creating a synthetic wound: (a-e) Stepwise fusion of tissue layers, (f) skin layer after removing the original wound via image completion, (g) original wound image, (h) final synthesis result in an advanced stage of disease. [52]

5.3 Generating Synthetic Wound Tissue

To generate a tissue layer for the use in wound synthesis described in Section 5.2 a large patch from a suitable image containing the respective tissue type is extracted. Based on that patch, a large texture of size 1024x768 is generated for the specific tissue type using example-based texture synthesis and blending following ideas of [14, 30, 59]. If there is no proper reference patch found (the matching distance is beyond a limit), which is mostly the case if the available amount of existing seed-texture is insufficient, similar patches from the patch database described in 4.6 were used for texture synthesis. Since these methods are time consuming several such textures for different tissue types have already been pre-calculated and stored for further use. The simulation software uses these textures at runtime. The lowest layer, the skin layer, is generated in a different way. The next Section 5.4 describes the procedure.

5.4 Completion of Wound Areas with Healthy Skin

The final result of a wound healing process is intact skin. When synthetically creating intact skin we have to find a way to remove the wound and its border (region) and replace the missing image parts properly by intact skin – preferably without big noticeable artefacts. As the author has already addressed this problem as part of his thesis in [52],

only a summary of the method is described here. The interested reader is referred to the mentioned paper for more information like details about the algorithm or performance comparisons with other papers.

In order to receive satisfying image completion results without wound border artefacts, it is necessary to completely remove the wound and its directly adjacent tissue, which mostly differs from intact skin in colour and texture. Removing the wound is a straight forward issue, since the wound mask is given. In case of DFS also a wound border region mask is needed for that purpose. The mask is created by wound border region detection, which we will see in Section 5.5.2. To remove even the last remnants of regions affected by the wound the given masks are morphologically dilated by a disk-shaped structuring element.

Remains the actual skin completion part. The implemented algorithm is based on circular patch based image completion. It works according to the steps depicted in Figure 5.10.




Figure 5.10: Concept of image completion technique: (a) section of original image; (b) clockwise filling order; rectangles show locations where a patch is fitted; the orange dash-dotted line represents the path of actual iteration; the white solid line represents the contour along which the image is cropped after each iteration; (c) shows blending step; (d) result after two iterations; (e) remaining gaps after last iteration, filled using the same blending technique as for the other patches; the rectangles depict the matching window; (f) final result. [52]

To start with, Figure 5.10(a) shows an image section containing the wound, in particular case one caused by PAOD. After removing the region affected by the wound, a contour path (orange dash-dotted line) is calculated in a way, that all patches of a given size that are located on the path with respect to their centre pixel overlap by around 30% of its area with surrounded skin (see Figure 5.10(b)). The overlapping region is matched with skin region in a search space in a grid-like manner based on cross-correlation. During experiments we recognised that the best fitting patches were located nearby the wound (presumed that the wound border is removed as described). Therefore a search mask allowing us to limit and optimise the search space, was introduced. Of course, all non skin regions were excluded from the search space, too. After a proper patch is found, it is merged to the output image, accordingly to Figure 5.10(c), whereas the upper left patch depicts the rectangular site of the image to be filled (superimposed by a seam mask), and the upper right patch shows the best matching patch (superimposed by the inverted seam mask). Adding these two patches results in a final patch (depicted underneath the two patches) that can be directly inserted into the output image. The white solid line represents the contour along which the image is cropped after each circular iteration. Figure 5.10(d) shows the intermediate result of the procedure after the second circular iteration. The remaining gaps after the last iteration as depicted in Figure 5.10(e) are handled using the same patch matching and blending steps as described before. The final result is depicted in 5.10f. The image completion method has been applied to various wound images and the results have been reviewed within the team and by a physician. Figure 5.11 shows some results of wound image completion, each line representing another example. The proposed method works well on more or less planar regions and when wounds are cleaned before the image is taken (see first and second row or Figure 5.11) otherwise it may produce sub-optimal results like that shown in the third row of Figure 5.11. The image completion method was accepted for the use in the final application by the consortium.

5.5 DFS Wound Border Treatment

In our images, the wound border region of DFS differs significantly from that of UCV and PAOD as already stated in 5.2.1. Most DFS wounds consist of a bright, white to yellowish region of varying width (see Figure 1.2). Such a border region is caused by partially soaked callused skin. If this region would not be taken into account in the synthetic replacement of the wound with healthy skin, unrealistic artefacts would be the result. For this reason, it has to be detected and removed before conducting wound synthesis operations.



Figure 5.11: Results of wound image completion based on three examples. a,e,i) show the original images. b,f,j) depict the image completion results. c,g,k) show the relevant section of original image; d,h,l) show the same section with completed skin.

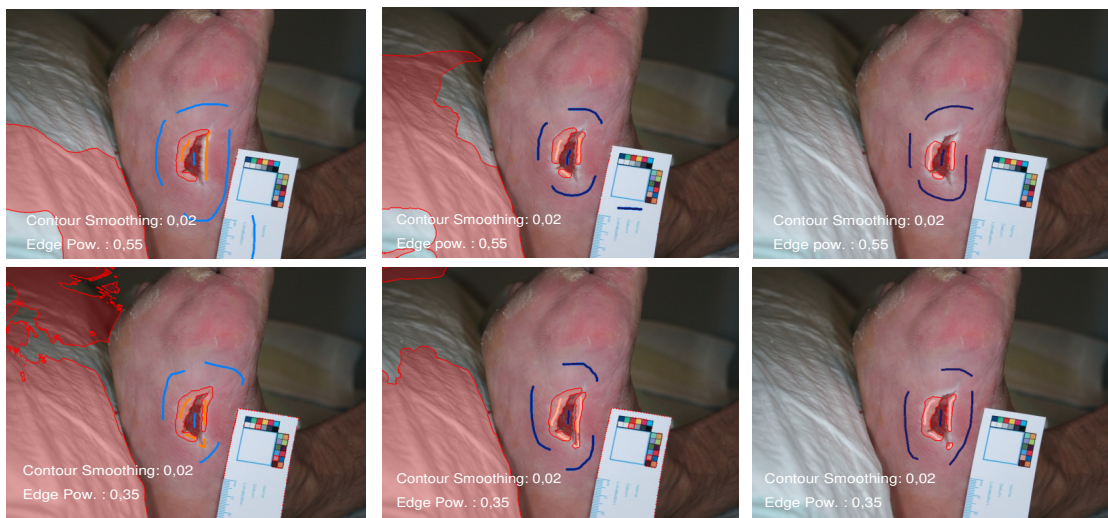
In the course of the search for a suitable method for segmenting this wound border region two interesting methods were short-listed and tested for their fitness in our application. The first was Interactive Total Variation Based Image Segmentation (TVSeg) introduced by Unger et al. in [57]. The interactive tool provided promising segmentation results based on sample images with smooth transitions between object and background. The second method was SRM, as it has provided excellent segmentation results on wound images so far (see sections 4.7.2 and 5.1.2).

5.5.1 Wound Border Region Segmentation Using TVSeg

According to their authors TVSeg [57] is based on a binary labelling into foreground and background using fast Total Variation minimisation of the Geodesic Active Contour energy with local constraints. In an optional second step TVSeg performs alpha-matting along the border of the binary segmentation.

As TVSeg inherently relies on the edge image derived from the input image, different edge related parameters affecting the final segmentation result have been exhaustively tested to achieve the best segmentation result for our use case. The *edge scale* value for

example controls the general influence of edges onto the segmentation result, whereas the *edge power* value controls the influence of the edge strength. If high, only strong edges influence the segmentation. Furthermore, the calculated edge images have been slightly smoothed in an edge preserving de-noising preprocessing step for better segmentation results. Let us now examine the results produced by the TVSeg when applied to our DFS images. Figure 5.12 and Figure 5.13 show the segmentation result of two different DFS input images. The left column of each figure shows the interactive segmentation results of the wound border region using standard constraints, the column in the middle using mixed constraints and the right column shows the results based on hard constraints offered by different brushes in the TVSeg GUI. The interactively drawn pink/orange brush strokes indicate foreground and the blue ones denote background constraints. The upper and lower rows differ in the used edge image related segmentation parameters. The specific values used are displayed in each result image.



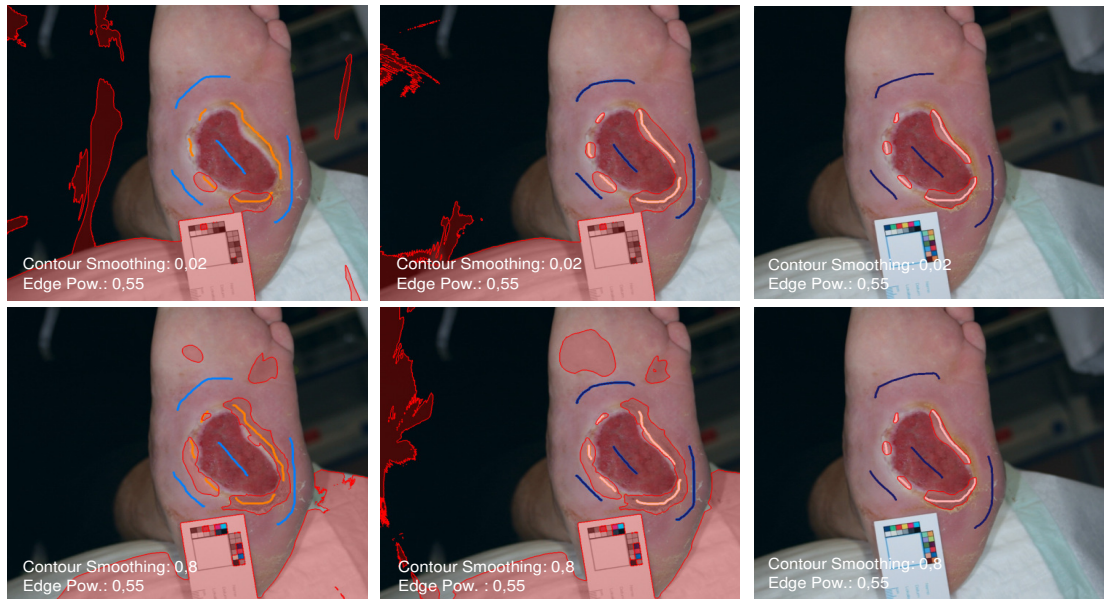
(a) Wound border Segmentation example using standard constraints.

(b) Wound border Segmentation example using mixed constraints.

(c) Wound border Segmentation example using hard constraints.

Figure 5.12: Examples of wound border region segmentation based on interactive TVSeg [57] using different parameters for edge detection. Interactively drawn pink/orange brush strokes indicate foreground and the blue ones background constraints.

Unfortunately, none of the tested parameters provided convincing segmentation results. Although several start regions for the foreground and background were defined interactively, no coherent or complete segmentation of the wound border region could be



(a) Wound border segmentation examples using standard constraints.

(b) Wound border segmentation examples using mixed constraints.

(c) Wound border segmentation examples using hard constraints.

Figure 5.13: Examples of wound border region segmentation based on interactive TVSeg [57] using different parameters for edge detection. Interactively drawn pink/orange brush strokes indicate foreground and the blue ones background constraints.

achieved. Also, an attempt, not shown in the two figures, of segmenting the wound border region together with the wound itself, brought no improvements. TVSeg is a powerful tool for segmenting objects from a background, of which the author was able to convince himself by means of several tests on non-wound images, but it is not right tool for separating the wound border region in our case of application.

5.5.2 Wound Border Region Segmentation Using SRM

Let us go through the second short-listed segmentation method based on SRM. In a first step a binary wound border region mask M_{min} was created by dilating the given wound mask of a DFS input image using a disk shaped structuring element of radius r_{min} . This mask represents the minimum boundary region for a DFS wound. In a second step the SRM algorithm was applied to the colour normalised input image, using a statistical complexity parameter $Q = 2^9$, which provided the best trade-off between low complexity and a fine segmentation. In the next step those segmented regions were selected, that

adjoin the wound border but do not exceed a specified maximum normal distance d_{max} from that border. This region set builds the binary wound border mask M_{seg} . In a last step the final wound border mask is generated $M_{seg} = OR(M_{min}, M_{seg})$. If necessary M_{seg} is cropped by the skin mask developed in Section 5.1.2. We have tested the method using all available DFS wounds and reviewed the results together with a physician and found the segmentation procedure suitable for our application. Figure 5.14 shows some results.



Figure 5.14: Examples of wound border region segmentation based on SRM [40].


5.5.3 Creation of DFS Wound Border

We have seen that we can detect the wound border environment of DFS. Because the DFS wound border region is many times more complex than that of PAOD and UCV, the wound deterioration simulation is not an easy task. Many attempts concerning this matter have failed. For this reason, the team decided in intensive discussions to extract reference wound border regions from all available images and store them in a database for access when needed. A wound edge region taken from the database is then adapted to the target wound contour by simple Euclidean transformations and merged as uppermost layer to the wound in a blending step using a proper wound border region mask. After reviewing the results in the team together with a wound expert, the method was approved.

5.6 Wound Cavity Synthesis

As already mentioned, the clinical picture of DFS also includes cavities that may occur at a more advanced stage of the disease and are an indication of the destruction of deeper tissue layers. We have seen in 4.7.2 how a cavity is detected in order to exclude it from the colour based tissue segmentation. In case of wound deterioration, these cavities must be artificially created.

In close consultation with the physician different approaches were tested. The best tested approach to create an artificial cavity, that came very close to a real one, was based on the euclidean distance transformation [18]. The transformation was applied on the wound contour mask in order to generate an attenuation filter. Directly applied to the wound the filter generates a nice cavity effect. The severity of generated cavity can be controlled by merely one blending parameter $\alpha_{intensity}$. The approach was tested using all available DFS images. The results have been reviewed together with a physician and found suitable. Figure 5.15 shows two examples of synthetically generated cavities using $\alpha_{int} = 0.5$. The left column depicts the original wound image. The column in the middle shows the synthetic cavity and the right column depicts the relevant zoomed image section.



(a) Original DFS wound images.

(b) Synthetically generated cavity.

(c) Zoomed image sections containing synthetic cavity.

Figure 5.15: Examples of synthetically generated cavities in DFS wounds.

Chapter 6

Wound Development Simulation

Contents

6.1	Available Image Data	65
6.2	Wound Development Simulation Chain	65
6.3	Simulation Results	70

6.1 Available Image Data

Before we devote ourselves to the wound development simulation, let us take a look at the available image data gathered until the end of project SimuWound. Table 6.1 provides information about available images of DFS in general and to assigned wound phases in particular. Furthermore, the table distinguishes between images without markers, those with a pure geometric marker and those containing a colour marker. As one can see, there exist quite many images in total, but very few to cover a complete wound development of a patient. In addition to this unpleasant aspect, there exist relatively few images containing a marker that allows sizing and colour normalisation of the wound. If we look at the data for PAOD in Table 6.2 and for UCV in Table 6.3, a similar picture emerges, although for the latter diagnosis the data situation looks a little better.

6.2 Wound Development Simulation Chain

Due to the poor data situation, it was necessary to refrain from deploying data-intensive machine-learning based techniques such as for example deep learning [35, 51] when generat-

Patient ID	Wound ID	Localisation	Total Images	Images containing marker					Images containing colour marker				
				Phase 1	Phase 2	Phase 3	Phase 4	Total	Phase 1	Phase 2	Phase 3	Phase 4	Total
Patient 01	01	Dig. II ped. dext.	3	x		x	x	3	x		x	x	3
Patient 02	02	Plant. dext.	5	x	x			4	x	x			4
Patient 03	03	Hall. phal. dist. plant. dext.	3			x	x	3			x	x	3
Patient 04	04	Pes lat. dext.	1				x	1				x	1
Patient 05	05	Hall. phal. dist. plant. sin.	2			x		2			x		2
Patient 06	06	Plant. dext.	1		x			1					-
Patient 07	07	Hall. phal. dist. plant. dext.	5	x	x	x	x	5	x	x	x	x	5
Patient 08	08	Plant. sin.	10	x	x	x		9	x	x	x		9
	09	Plant. dext.	1	x				1	x				1
Patient 09	10	Plant. dext.	4		x			2		x			2
Patient 10	11	Plant. sin.	4	x				4	x				4
Patient 11	12	Plant. dext.	10	x	x	x		8	x	x	x		7
Patient 12	14	Plant. sin.	6	x		x		5	x		x		5
Patient 13	15	Dig. I/II ped. dors. dext.	2		x			2		x			2
Patient 14	16	Plant. dext.	5	x		x	x	5	x		x	x	5
Patient 15	17	Plant. dext.	14	x		x	x	4					-
Patient 16	18	Plant. dext.	7	x	x	x		6	x	x	x		5
Patient 17	19	Plant. dext.	4	x	x			4	x	x			4
Patient 18	20	Plant. sin.	2		x			2		x			2
Patient 19	21	Dors. dext.	6		x	x	x	6		x	x	x	6
Patient 20	22	Dors. dext.	5	x		x		5	x		x		5
Not attributable	00	-	27	x	x	x		22	x	x	x		22
Total			127					104					97
Total attributable to patients			100					82					75

Table 6.1: Available image data for diagnosis DFS.

Patient ID	Wound ID	Localisation	Total Images	Images containing marker					Images containing colour marker				
				Phase 1	Phase 2	Phase 3	Phase 4	Total	Phase 1	Phase 2	Phase 3	Phase 4	Total
Patient 01	01	Crus. fib. sin.	5	x	x	x		5	x	x	x		4
	02	Crus. tib. sin.	3	x		x		3	x		x		3
Patient 02	03	Crus. tib.post.sin.	4			x	x	4			x	x	3
Patient 03	04	Calc. med. dext.	13	x	x	x	x	10					-
Patient 04	05	Crus. sin./dext.?	1				x	1				x	1
	06	Crus. fib. sin.	22		x	x	x	18			x	x	9
Patient 05	07	Pes. dors. sin.	29		x	x		22		x	x		9
Patient 06	08	Hall. med.dext.	13				x	4					-
	09	Mall. lat. dext.	7		x	x		5			x		1
	10	Mall. lat. sin.	9			x	x	5			x		1
Patient 07	11	Mall. med. dext.	8		x	x	x	6			x	x	2
	12	Mall. med. sin.	9			x	x	6			x		1
Patient 08	13	Crus. fib. dext.	7		x	x		7		x	x		7
	14	Mall. tib. sen.	7			x		5			x		5
Patient 09	15	Mall. tib. dext.	6			x		5			x		5
Patient 10	16	Mall. fib. dext.	6	x	x	x		6	x	x			4
Patient 11	17	Crus. tib. dext.	4		x	x		4		x	x		4
Patient 12	18	Crus. fib. dext.	10		x	x	x	10		x	x	x	10
Not attributable	00	-	34	x	x	x	x	24	x	x	x	x	24
Total			197					150					93
Total attributable to patients			163					126					69

Table 6.2: Available image data for diagnosis PAOD.

ing a suitable wound simulation model. The finally implemented model for the simulation of wounds is based on the evaluation of a few existing wound developments with regard to wound size and tissue type proportions, furthermore, on the work of [61] addressing healing times concerning DFS based on a study, and mainly on the experience of our physician, a certified wound manager. The model parameters are stored in a structure and can be easily adjusted as new specifications or new (more) ground-truth data is available or a need for adjustment arises on part of physicians.

As already explained in Section 5.2 a synthetic wound is essentially generated by

Patient ID	Wound ID	Localisation	Total Images	Images containing marker					Images containing colour marker				
				Phase 1	Phase 2	Phase 3	Phase 4	Total	Phase 1	Phase 2	Phase 3	Phase 4	Total
Patient 01	01	Mall. tib. sin.	7		x	x		7		x	x		7
Patient 02	02	Crus. fib. sin.	2	x				1					-
	03	Crus. tib. sin.	45		x	x		28		x			2
Patient 03	04	Mall. fib. sin.	1			x		1			x		1
	05	Mall. tib. sin.	1			x		1			x		1
Patient 04	06	Mall. fib. dext.	1		x			1		x			1
	07	Mall. fib. sin.	1		x			1		x			1
	08	Mall. tib. sin.	5		x	x		5		x	x		5
Patient 05	09	Mall. ?	1			x		1			x		1
Patient 06	10	Mall. tib. dext.	5	x	x		x	5	x	x		x	5
Patient 07	11	Crus. fib.dext.	10		x	x		7		x			1
Patient 08	12	Crus. fib.sin.	5			x	x	5			x	x	3
Patient 09	13	Mall. tib. sin.	5		x	x		5		x	x		5
Patient 10	14	Crus. fib.sin.	5			x	x	5			x	x	4
	15	Crus. tib. dext.	2				x	2				x	2
	16	Crus. tib. sin.	2				x	2				x	2
Patient 11	17	Crus. fib. dext.	4		x	x	x	4		x		x	3
Patient 12	18	Crus. fib. sin.	1				x	1				x	1
Patient 13	19	Crus. dext.	10	x	x	x	x	10	x	x			2
Patient 14	20	Mall. tib. sin.	7	x	x			7	x	x			7
Patient 15	21	Crus. tib. sin.	17	x	x	x	x	16	x	x	x	x	12
	22	Crus. tib. dext.	8		x	x	x	8		x	x	x	7
Patient 16	23	Crus. tib.sin.	14		x	x		14		x	x		14
Patient 17	24	Mall. fib. sin.	8	x	x	x		8	x	x	x		8
	25	Mall. tib. sin.	5	x	x			5	x	x			4
	26	Mall. fib. dext.	3	x				3	x				3
Patient 18	27	Mall. med. sin.	9		x	x		8					-
Patient 19	28	Crus. tib. dext.	2				x	2				x	2
Patient 20	29	Crus. fib.sin.	19	x	x	x	x	14					-
Patient 21	30	Crus. tib.sin.	1			x		1			x		1
Patient 22	31	Mall. fib. dext.	1		x			1					-
Not attributable	00	-	51	x	x	x	x	48	x	x	x	x	48
Total			258					227					153
Total attributable to patients			207					179					105

Table 6.3: Available image data for diagnosis UCV.

blending different synthetic tissue layers. Each tissue layer is associated with a filtered mask that is altered by a combination of scaling steps, morphological operations, and threshold methods to simulate wound healing or deterioration according to the mentioned model parameter structure. Parameters that are necessary for wound cavity synthesis in case of DFS and additional effect parameters for inflammation and swelling are also part of the structure.

Inflammations were simulated by filtering a wound mask using a huge Gaussian kernel. The so generated filter mask was applied to the skin layer mainly strengthen the red RGB channel. Note that this operation was only applied to the surrounding skin and not to the wound itself. The inflammation effect can be observed in the wound deterioration simulation for UCV shown in Figure 6.6 and for those of DFS in Figure 6.4. Finally, the swelling of the foot of a DFS-affected patient was simulated using a distortion field generated by Thin Plate Spline (TPS) [7] according to specifications of a physician. Figure 6.4 shows the swelling as a concomitant effect of wound deterioration in DFS. Both effects have been tested on various images and found to be realistic by physicians.

In all simulation steps great importance was placed on randomness. Thus, for most parameters concerning wound growth and healing, small Gaussian noise with a defined

variance was added. This procedure ensures that one and the same input image with the same given personal circumstances lead to similar but not identical simulations.

The necessary preprocessing steps for the wound evolution simulation chain are explained in a state flow diagram shown in Figure 6.1 and the actual wound evolution simulation process itself is explained in a second state flow diagram depicted in Figure 6.2. The two state flow diagrams combine all marker and wound related analysis, synthesis and simulation steps that were presented in this thesis so far. So there all threads come together.

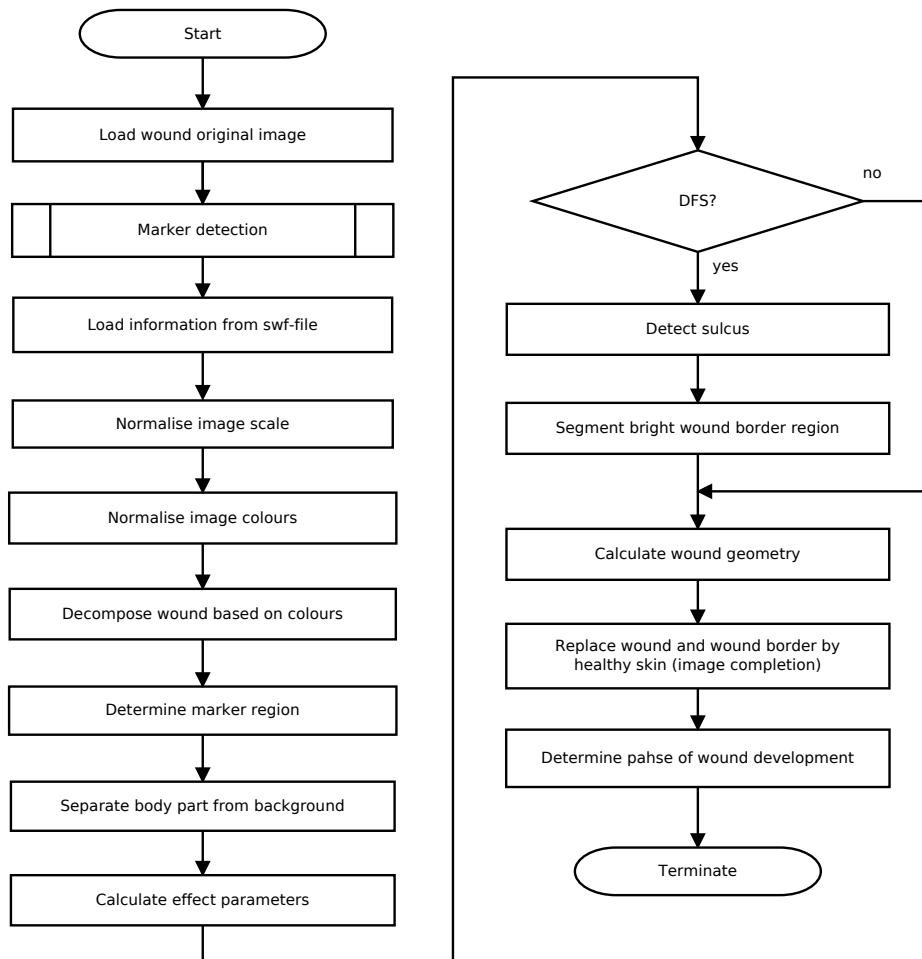


Figure 6.1: State flow diagram of image preprocessing.

6.3 Simulation Results

Finally, some wound development simulations are presented here. Figure 6.3 shows wound healing for DFS and Figure 6.4 shows the wound deterioration for DFS. Note the additional swelling and inflammation effect. As one can see it forms a cavity starting in Figure 6.4(d).




Figure 6.3: Simulation of healing for DFS. a) Original image of DFS at begin of treatment., b-f) predicted stages of disease with a time delay of five days between each stage.




Figure 6.4: Simulation of deterioration for DFS. a) Original image of DFS wound at an early stage, b-f) predicted stages of disease with a time delay of one week between each stage, if not treated.

Finally Figure 6.5 presents a healing simulation for UCV and Figure 6.6 synthetic deterioration. Since the simulation path of PAOD is similar to that of UCV, the figure is omitted here.




Figure 6.5: Simulation of healing for UCV. a) Original image of leg ulcer at begin of treatment, b-f) predicted healing with a time delay of one week between each stage.



Figure 6.6: Simulation of deterioration for UCV. a) Original image of leg ulcer at an early stage, b-f) predicted stages of disease with a time delay of two weeks between each stage.

Chapter 7

Conclusion and Outlook

7.1 Summary of Contribution

The main contributions of this thesis have been the analysis and description of visual attributes of a chronic wound with respect to a related diagnosis and the generation of synthetic wounds based on the extracted features. Furthermore, a possible way to simulate different wound evolutions depending on given personal circumstances was outlined. The outcome of this work was a substantial contribution to the successful development of a wound simulation tool within the research project SimuWound, which was approved for everyday clinical use.

In more detail, chapter 1 addressed the actual healthcare situation concerning chronic wounds and described the reasons for the development of the wound simulation tool and formulated the clear goal of patient motivation to actively participate in the therapy in order to accelerate the healing process and drastically reduce healthcare costs. After defining the scope of work a short description of the covered diagnoses DFS, PAOD and UCV was given.

Chapter 2 provided an overview of existing work in the field of image based wound analysis and wound development prediction.

Chapter 3 presented a well chosen, reliable and easy to use image acquisition system that enables the capturing of high quality wound images. Furthermore, a carefully developed detection algorithm for different size and colour reference markers, achieving high detection rates in conducted evaluation experiments was introduced.

Chapter 4 addressed the analysis of wounds. The consensually classification of wounds into phases of a disease within the project team as well as definition of wound features

relevant for image processing were explained. Moreover, the derivation of promising state of the art wound descriptors from the defined features was presented in detail. Based on the descriptors a decision tree method for automated wound phase classification that even works with few image training data was developed and evaluated. A suitable method for finding similar wound patches in a database for later use in wound synthesis were analysed and exhaustively tested by the author as well. In addition, methods for detecting cavities in a wound were implemented and evaluated and finally, a suitable approach of segmenting a wound into several tissues was presented.

Chapter 5 tackled wound synthesis. In order to solve the skin segmentation task in wound images, an existing approach and a newly developed one were tested and evaluated in detail. Furthermore, a new layer concept for generating a synthetic wound was presented and all necessary techniques and steps were explained in detail. These included a newly developed image completion method optimised for wounds, adapting existing texture synthesis methods for wound tissue layer generation and a newly developed method for wound border region detection in case of DFS. Finally, a newly method for the generation of synthetic but realistic looking wound cavities was presented.

Chapter 6 was about wound evolution simulation. Starting with, an overview of available image data in the project was given and the simulation of healing and worsening of wounds was described. Moreover, some additional effects addressing inflammation and swelling were introduced and detailed flow chart diagrams explaining the work-flow of the developed wound simulation tool were given. Finally, some simulation results for different diagnoses regarding healing and worsening were presented.

7.2 Conclusion, Discussion and Outlook

The wound simulation tool and the deployed methods presented in this thesis worked well and proved to be applicable for daily clinical routine. The tests showed that even non-expert users (clinical personal) upon a short training session can use the system and first results showed the acceptance by patients. The simulation tool also has the potential for the use in mobile devices or as valuable tool in the training of prospective physicians and other persons involved in wound care in the clinical but also in the private sector.

Nevertheless, there are still possibilities to further enhance the system, respectively the appearance of simulated wounds to motivate patients even more in active wound treatment. For example, the applied methods for the analysis of wound parameters and their use in

wound simulation could partly be replaced by powerful deep learning approaches [35, 51]. However, the actual bad data situation already mentioned in the work does not permit such data-intensive methods at the moment. But the wound analysis and synthesis developed in this work have the potential to provide suitable data for training of deep learning methods. There is also room for improvement in the automation of manual wound border segmentation, to relieve the physician. Automated approaches for wound border segmentation already exist (see for example [15, 49]). Whether these approaches are capable to deal with the great variability in all presented chronic wounds due to different wound margins but not least because of the high variability of wound border regions and surrounding skin is open. Automated skin detection could also be predestined for deep learning methods just mentioned. With sufficient data, the perceivable variability in manual wound segmentation observed during the project between individual skilled personnel could be addressed with these methods, too.

As already described in the introduction to this work, the presented method of wound analysis and synthesis is limited to more or less planar skin areas due to the two-dimensional nature of a wound image. The use of a 3D camera system like that of EKARE [15] and appropriate adaptation of the algorithms to the three-dimensional conditions would considerably expand the field of application of the presented simulation tool. In addition to the planar skin areas then finer limbs such as fingers or toes could be taken into account or wounds that have manifested around a limb. The visual impression of a wound simulation, which is important for increasing compliance, could, according to the author, be significantly improved by 3D methods. Also the depth of the wound, which is a parameter used to assess the severity of an injured skin area, can be determined to a certain extent by 3D methods, as also evidenced by EKARE. This is not possible with 2D methods or only insufficiently and indirectly using workarounds (see cavity detection in chapter 4.7.2). The depth parameter could provide a valuable contribution to the improvement of wound classification and thus also to wound simulation.

The set of wound parameters used for the wound analysis can be extended even further. This work deals exclusively with RGB colour images in the visible light spectrum. The use of thermal imaging or multi-spectral cameras can provide valuable insights into processes beneath the visible wound surface. Inflammatory reactions in the wound or differently perfused wound areas are detectable in the near infra-red spectrum (see [27] and [31]) and thus can provide interesting additional information regarding the disease state. This information could be used to improve the presented wound analysis methods. Furthermore,

there is a possibility to enhance the simulation tool concerning the covered diagnoses which were limited to DFS, PAOD and UCV in project SimuWound and in this thesis. An extension to pressure ulcers, burn wounds and other diseases associated with skin injuries would significantly expand the range of application of the simulation tool.

The outlined improvement potential can serve as the basis for planned further developments of the simulation tool. In any case, a major focus of future work will have to be on the generation of more training image data. The actual positive effect of the wound simulation tool on the motivation of the patient and the shortening of treatment duration will probably only be assessed in a long-term study with extensive statistical evaluations.

Appendix A

List of Publications

My work at JOANNEUM RESEARCH led to the following peer-reviewed publication:

Image Completion Optimised for Realistic Simulations of Wound Development

Michael Schneeberger, Martina Uray and Heinz Mayer

In: Pattern Recognition. Proceedings of Joint DAGM and OAGM Symposium

Lecture Notes in Computer Science, volume 7476

August 2012, Graz, Austria

Springer, Berlin Heidelberg

pages 448–457.

Abstract: Treatment costs for chronic wound healing disturbances have a strong impact on the health care system. In order to motivate patients and thus reduce treatment times there was the need to visualize possible wound developments based on the current situation of the affected body part. Known disease patterns were used to build a model for simulating the healing as well as the worsening process. The key point for the construction of possible wound stages was the creation of a nicely fitting texture including all representative tissue types. Since wounds are mostly circularly shaped, as first step of the healing an image completion based on radial texture synthesis of small patches from the healthy tissue surrounding the wound was developed. The radial information of the wound border was used to optimize the overlap between individual patches. In a similar way complete layers of all other appearing tissue types were constructed and superimposed using masks representing trained possible appearances. Results show that the developed texture synthesis together with the trained knowledge is perfectly suited to construct realistic wound images for different stages of the disease.

Bibliography

- [1] Ahonen, T., Hadid, A., and Pietikäinen, M. (2006). Face description with local binary patterns: Application to face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(12):2037–2041.
- [2] Akestes GmbH (2018). WundManager. Website. <http://www.akestes.de/>, accessed July 20, 2018.
- [3] American Colleg of Cardiology Foundation, American Heart Association, Society for Cardiovascular Angiography and Interventions, Society of Interventional Radiology, Society for Vascular Medicine, and Society for Vascular Surgery (2011). *Management of Patients With Peripheral Artery Disease*. Elsevier, New York, USA.
- [4] Arbeitsgemeinschaft der Wissenschaftlichen Medizinischen Fachgesellschaften (AWMF) (2011). *Management of Patients With Peripheral Artery Disease*. Elsevier, New York, USA.
- [5] Barnes, C., Shechtman, E., Goldman, D. B., and Finkelstein, A. (2010). The Generalized PatchMatch Correspondence Algorithm. In *Proc. of 11th European Conference on Computer Vision: Part III*, pages 29–43, Berlin, Heidelberg. Springer-Verlag.
- [6] Bon, F.-X., Briand, E., Guichard, S., Couturaud, B., Revol, M., Servant, J.-M., and Dubertret, L. (2000). Quantitative and kinetic evolution of wound healing through image analysis. *IEEE Transactions on Medical Imaging*, 19(7):767–772.
- [7] Bookstein, F. L. (1989). Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585.
- [8] Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1984). *Classification and regression trees*. Brooks/Cole Publishing, Monterey, CA, USA.
- [9] Brunelli, R. (2009). *Template Matching Techniques in Computer Vision: Theory and Practice*. John Wiley & Sons, Chichester, UK.

-
- [10] Calonder, M., Lepetit, V., Ozuysal, M., Trzcinski, T., Strecha, C., and Fua, P. (2012). Brief: Computing a local binary descriptor very fast. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(7):1281–1298.
- [11] Calonder, M., Lepetit, V., Strecha, C., and Fua, P. (2010). Brief: Binary robust independent elementary features. In Daniilidis, K., Maragos, P., and Paragios, N., editors, *Proceedings of 11th European Conference on Computer Vision (ECCV 2010)*, volume 6314 of *Lecture Notes in Computer Science*, pages 778–792. Springer Berlin / Heidelberg.
- [12] Chang, C.-C. and Lin, C.-J. (2011). LIBSVM: A library for support vector machines. *ACM Trans. on Intelligent Systems and Technology*, 2:27:1–27:27.
- [13] Criminisi, A., Pérez, P., and Toyama, K. (2004). Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on Image Processing*, 13:1200–1212.
- [14] Efros, A. A. and Freeman, W. T. (2001). Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 341–346, New York, NY, USA. ACM.
- [15] EKARE Inc. (2018). inSight. Website. <https://ekare.ai/>, accessed July 20, 2018.
- [16] Finlayson, G. D., Hordley, S. D., and Hubel, P. M. (2001). Color by correlation: a simple, unifying framework for color constancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1209–1221.
- [17] Gallenkemper, G. (2008). Guidelines for diagnosis and therapy of venous ulcers (version 8 2008). *Phlebologie*, 37(6):308–329.
- [18] Gonzalez, R. C. and Woods, R. E. (2008). *Digital image processing*. Prentice Hall, Upper Saddle River, N.J.
- [19] Guo, J., Zhao, Y., and Cai, A. (2010). A reliable method for paper currency recognition based on lbp. In *2010 2nd IEEE International Conference on Network Infrastructure and Digital Content*, pages 359–363.

- [20] Harris, C. and Stephens, M. (1988). A combined corner and edge detector. In *Proceedings of the 4th Alvey Vision Conference*, pages 147–151.
- [21] Hartley, R. I. and Zisserman, A. (2004). *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition.
- [22] Hsu, C.-W., Chang, C.-C., and Lin, C.-J. (2003-2016). A practical guide to support vector classification. Technical report, Department of Computer Science, National Taiwan University.
- [23] International Diabetes Federation (2013/14). *IDF Diabetes Atlas*. International Diabetes Federation, Brussels, Belgium, 6th edition. 2014 update.
- [24] International Working Group on the Diabetic Foot (2011). *International Consensus on the Diabetic Foot and Practical Guidelines on the Management and Prevention of the Diabetic Foot*. IWGDF, Amsterdam, the Netherlands.
- [25] Jalomed GmbH (2012). JalomedWD. Website. <http://www.jalomed.de/de/>, accessed July 01, 2018.
- [26] Jones, M. J. and Rehg, J. M. (2002). Statistical color models with application to skin detection. *International Journal of Computer Vision*, 46(1):81–96.
- [27] KENT Imaging (2018). KD203. Website. <https://ekare.ai/>, accessed July 20, 2018.
- [28] Kolesnik, M. and Fexa, A. (2005). Multi-dimensional color histograms for segmentation of wounds in images. *Lecture Notes in Computer Science*, 3656:1014–1022.
- [29] Komodakis, N. and Tziritas, G. (2007). Image completion using efficient belief propagation via priority scheduling and dynamic pruning. *Image Processing*, 16(11):2649–2661.
- [30] Kwatra, V., Essa, I., Bobick, A., and Kwatra, N. (2005). Texture optimization for example-based synthesis. *ACM Transactions on Graphics*, 24(3):795–802.

-
- [31] Landsman, A. (2018). Near infrared spectroscopy & predicting the likelihood of future wound healing. *Today's Wound Clinic*, 12(1):1452–1458.
- [32] Lange, D. U. (2008). *Potenziale und Grenzen der digitalen Wundfotografie in der Klinik und telematischen Nachsorge von Patienten mit Wunden*. PhD thesis, Ruhr-Universität Bochum.
- [33] Lashkia, G. V. and Anthony, L. (2003). An inductive learning method for medical diagnosis. *Pattern Recognition Letters*, 24(1-3):273–282.
- [34] Lawall, H. (2017). Diabetisches Fuß-Syndrom über Polyneuropathie, Durchblutungsstörungen und Zweitmeinungen. In *Deutscher Gesundheitsbericht Diabetes 2017: Die Bestandsaufnahme*, pages 83–95, Mainz, Germany. Deutsche Diabetes Gesellschaft (DDG) und Deutsche Diabetes-Hilfe (diabetesDE), Kirchheim + Co GmbH.
- [35] LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521:436–44.
- [36] Matas, J., Chum, O., Urban, M., and Pajdla, T. (2002). Robust wide baseline stereo from maximally stable extremal regions. In Rosin, P. L. and Marshall, A. D., editors, *Proceedings of 13th British Machine Vision Conference 2002 (BMVC 2002)*, pages 384–393, Cardiff, UK. British Machine Vision Association.
- [37] MATLAB (2017). *Version 9.2.0 (R2017a)*. The MathWorks, Natick, Massachusetts, USA.
- [38] Michael J. Swain, D. H. B. (1991). Color indexing. *International Journal of Computer Vision*, 7(1):11–32.
- [39] Mukherjee, R., Manohar, D. D., Das, D. K., Achar, A., Mitra, A., and Chakraborty, C. (2014). Automated tissue classification framework for reproducible chronic wound assessment. *BioMed Research International*, 2014:9 pages.
- [40] Nock, R. and Nielsen, F. (2004). Statistical region merging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(11):1452–1458.

- [41] Ojala, T., Pietikäinen, M., and Harwood, D. (1994). Performance evaluation of texture measures with classification based on kullback discrimination of distributions. In *Proceedings of 12th International Conference on Pattern Recognition*, volume 1, pages 582–585.
- [42] Ojala, T., Pietikäinen, M., and Mäenpää, T. (2001). A generalized local binary pattern operator for multiresolution gray scale and rotation invariant texture classification. In Singh, S., Murshed, N., and Kropatsch, W., editors, *Proceedings of Second International Conference on Advances in Pattern Recognition (ICAPR 2001)*, pages 397–406, Rio de Janeiro, Brazil. International Association for Pattern Recognition (IAPR), Springer, Berlin.
- [43] Ojala, T., Pietikäinen, M., and Mäenpää, T. (2002). Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):971–987.
- [44] Paletta, L. and Uray, M. (2011). Technischer Report Prognosemodul - Clusteranalyse im Merkmalsraum der Wundenbeschreibung. Technical report, Machine Vision Applications Research Group, Institute DIGITAL, JOANNEUM RESEARCH.
- [45] Peate, I. and Glencross, W. (2011). *Wound Care at a Glance*. John Wiley and Sons, Chichester, UK.
- [46] Pietikäinen, M., Ojala, T., and Xu, Z. (2000). Rotation-invariant texture classification using feature distributions. *Pattern Recognition*, 33(1):43–52.
- [47] Pietikäinen, M. and Zhao, G. (2016). Two decades of local binary patterns: A survey. *CoRR*, abs/1612.06795.
- [48] Prodan, A., Campean, R., Revnic, C., and Prodan, R. (2009). Strategies for wound image understanding. In *Proceedings of 7th ACS/IEEE International Conference on Computer Systems and Applications (AICCSA 2009)*, pages 1018–1024, Rabat, Morocco. IEEE Computer Society, IEEE, Los Alamitos, USA.

- [49] Prodan, A., Campean, R., Rusu, M., Revnic, C., Mitrea, P., and Prodan, R. (2010). Toward a model for wound healing simulation. In *2010 Fourth UKSim European Symposium on Computer Modeling and Simulation*, pages 115–120.
- [50] Saxen, F. and Al-Hamadi, A. (2014). Color-based skin segmentation: An evaluation of the state of the art. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 4467–4471.
- [51] Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117.
- [52] Schneeberger, M., Uray, M., and Mayer, H. (2012). Image completion optimised for realistic simulations of wound development. In Pinz, A., Pock, T., Bischof, H., and Leberl, F., editors, *Pattern Recognition*, volume 7476 of *Lecture Notes in Computer Science*, pages 448–457. Springer, Berlin Heidelberg.
- [53] Seevinck, J., Scerbo, M. W., Belfore, L. A., Weireter, L. J., Crouch, J. R., Shen, Y., McKenzie, F. D., Garcia, H. M., Girtelschmid, S., Baydogan, E., and Schmidt, E. A. (2006). A simulation-based training system for surgical wound debridement. *Studies in Health Technology and Informatics*, 119:491–496.
- [54] Sony (2010). *Interchangeable Lens Digital Camera alpha 33 and 35: Instruction Manual*. Sony Corporation, Thailand.
- [55] The MathWorks (2010). *Image Processing Toolbox R2017a: Documentation*. The MathWorks, Natick, Massachusetts, USA.
- [56] Treuillet, S., Albouy, B., and Lucas, Y. (2009). Three-dimensional assessment of skin wounds using a standard digital camera. *MedImg*, 28(5):752–762.
- [57] Unger, M., Pock, T., Cremers, D., and Bischof, H. (2008). TVSeg - interactive total variation based image segmentation. In *British Machine Vision Conference 2008*, Leeds, UK.

- [58] van Poucke, S., vander Haeghen, Y., Vissers, K., Meert, T., and Jorens, P. (2010). Automatic colorimetric calibration of human wounds. *BMC Medical Imaging*, 10(7):1–11.
- [59] Wei, L.-Y., Lefebvre, S., Kwatra, V., and Turk, G. (2009). State of the art in example-based texture synthesis. In *Eurographics 2009, State of the Art Report, EG-STAR*. Eurographics Association.
- [60] Wild, T., Prinz, M., Fortner, N., Krois, W., and S. Stremitzer, K. S., and Hölzenbein, T. (2008). Digital measurement and analysis of wounds based on colour segmentation. *European Surgery*, 40(1):5–10.
- [61] Zimny, S. and Pfohl, M. (2005). Healing times and prediction of wound healing in neuropathic diabetic foot ulcers: a prospective study. *Experimental and clinical endocrinology and diabetes. Official journal of German Society of Endocrinology [and] German Diabetes Association*, 113:90–93.