

Politecnico di Torino

Politecnico di Torino

Master's Degree in Biomedical Engineering - Biomedical Instrumentation

Academic Year 2025/2026

Quantitative OCT and OCTA Imaging for the Analysis of Radiation Therapy Treatment on Skin Cancer Patients

Supervisors:

Prof. Kristen Mariko MEIBURGER
Prof. Mengyang LIU
Dr. Giulia ROTUNNO
M.Sc. Giulio CARMASSI

Candidate:

Vincenza Claudia DICUONZO

Acknowledgements

I am deeply grateful to my supervisors, Prof. Kristen Mariko Meiburger and Prof. Mengyang Liu, for giving me the opportunity to work on this project and for their guidance and support throughout this journey.

My heartfelt thanks also go to Giulio, Lisa and Giulia for their invaluable help and advice; their support was essential to the completion of this work.

Finally, I am grateful to my colleagues in the office and laboratory at the Medical University of Vienna for creating a great working atmosphere, which made working on this project a truly positive and engaging experience.

Abstract

Optical Coherence Tomography (OCT) and Optical Coherence Tomography Angiography (OCTA) are non-invasive imaging modalities capable of providing high-resolution structural and microvascular information from biological tissues. In dermatology, their ability to characterize tissue morphology and vascular architecture makes them promising tools for lesion assessment and longitudinal monitoring of treatment-induced changes.

This thesis focuses on patients affected by non-melanoma skin cancer (NMSC), including basal cell carcinoma, squamous cell carcinoma, and actinic keratosis, undergoing superficial radiotherapy. Although radiotherapy represents an effective treatment option, particularly for patients unsuitable for surgery, it induces structural and vascular changes in lesions and may also lead to radiation-related toxicity. Quantifying these alterations may provide useful information for lesion assessment and treatment-response monitoring.

The study was conducted on a longitudinal clinical dataset acquired at the Medical University of Vienna, comprising OCT and OCTA volumes collected before treatment, immediately after radiotherapy, and at 3 and 6-month follow-up examinations. Starting from an existing OCTA analysis framework, several methodological developments were introduced to improve robustness, reproducibility, and clinical applicability. Motion-artifact correction was implemented to improve OCTA image quality. In addition, an automated vascular segmentation approach was evaluated against the previously adopted manual procedure to assess whether comparable quantitative vascular information could be obtained while reducing processing time and operator dependency.

The original feature set comprised OCTA-derived vascular parameters and radiomic descriptors characterizing image intensity, texture, heterogeneity, and lesion morphology. OCT-derived structural features describing epidermal and dermal morphology were added to provide complementary information on tissue architecture. The resulting multimodal feature space integrated vascular, radiomic, and structural information to provide a more complete description of skin tissue and its evolution during treatment.

The quantitative biomarkers were investigated both for their ability to discriminate the acquisition timepoints and for their association with radiation-induced skin toxicity. Patient-aware statistical and machine-learning approaches were used to evaluate classification performance. The inclusion of OCT-derived information enhanced timepoint discrimination, achieving a macro-averaged multiclass Area Under Curve (AUC) of 0.78. Toxicity analysis showed high discrimination (AUC = 0.88) when baseline information was combined with early follow-up changes. The addition of OCT-derived features improved toxicity discrimination in most model configurations, with AUC increases up to 0.11. Moreover morphological OCT features were repeatedly retained during stability selection, suggesting complementary contribution of structural information.

Overall, this work establishes an automated and multimodal quantitative OCT/OCTA framework for investigating skin lesions and longitudinal tissue response to radiotherapy. By integrating vascular, radiomic, and structural biomarkers, the proposed methodology provides a foundation for objective and non-invasive tools that may support lesion characterization, treatment monitoring, toxicity assessment, and personalized clinical management of patients with NMSC.

Table of Contents

List of Figures VIII

1 Introduction 1
- 1.1 Clinical Background: Non-Melanoma Skin Cancer . . . 1
  - 1.1.1 Characterization of BCC, SCC, and AK . . . 1
  - 1.1.2 Radiotherapy Management and Treatment Alternatives . . . 2
  - 1.1.3 Radiation-Induced Skin Toxicity . . . 3
- 1.2 Optical Coherence Tomography (OCT) . . . 4
  - 1.2.1 Physical Principles and Signal Acquisition . . . 4
  - 1.2.2 OCT Techniques and Parameters . . . 6
  - 1.2.3 Morphological OCT vs. Functional OCT-Angiography (OCTA) 9
  - 1.2.4 State of the Art in OCT/OCTA Analysis . . . 9
- 1.3 Purpose of the Study . . . 11

2 Materials and Methods 13
- 2.1 Dataset and Clinical Protocol . . . 13
  - 2.1.1 Patient Cohort and pre-existing Framework . . . 13
  - 2.1.2 Previous processing state . . . 14
  - 2.1.3 Radiotherapy Protocol . . . 15
  - 2.1.4 Laboratory System . . . 15
  - 2.1.5 Scanning Protocol . . . 17
  - 2.1.6 Raw Data Acquisition and Reconstruction . . . 18
  - 2.1.7 Acquisition Challenges and Limitations . . . 19
- 2.2 Data Optimization Pipeline . . . 19
- 2.3 OCT Morphological Processing and Feature Extraction Pipeline . . 24
  - 2.3.1 Volume Preprocessing and Geometric Realignment . . . 26
  - 2.3.2 Skin Surface and Dermal-Epidermal Junction Segmentation 34
  - 2.3.3 Quantitative Feature Extraction . . . 42
- 2.4 Statistical Analysis . . . 44
  - 2.4.1 Longitudinal Classification . . . 44
  - 2.4.2 Clinical Toxicity Scoring . . . 45

2.4.3 Toxicity Discrimination Model 46
2.4.4 Performance Evaluation Metrics 48

3 Results 50
3.1 OCTA Data Optimization 50
3.1.1 Automatic Segmentation Outcomes 50
3.1.2 Impact of Motion Artifact Correction on Longitudinal Classification 52
3.2 OCT Morphological Analysis 56
3.2.1 OCT Segmentation and Morphological Maps 56
3.2.2 Distribution of OCT Morphological Features 61
3.3 Contribution of OCT Morphological Features to Quantitative Classification 65
3.3.1 Longitudinal Classification: OCTA/Radiomics vs. OCT/OCTA/Radiomics 65
3.3.2 Toxicity Classification and Contribution of OCT Morphological Features 67

4 Discussion and Conclusions 72
4.1 Data Optimization 72
4.2 Interpretation of OCT/OCTA Quantitative Findings 74
4.3 Multimodal Imaging and Radiation-Induced Toxicity Assessment . . 77
4.4 Limitations and Future Developments 78
4.5 Conclusions 80

List of Figures

1.1 Standard OCT system based on Michelson architecture. LS=low time-coherence light source, PC= personal computer [4] 5
1.2 Standard time-domain (TD) and Fourier-domain (FD) Michelson interferometer configurations used in OCT. [5] 7

2.1 Schematic representation of the laboratory SS-OCT system used for clinical acquisitions. Taken from [23] 16
2.2 SS-OCT system used for clinical acquisitions. Taken from [15] . . . 17
2.3 Overview of the data-optimization workflow. The analysis comprised two main stages: evaluation of automated vascular segmentation against the manual reference and assessment of the segmentation after motion-artifact correction of selected corrupted OCTA volumes. 20
2.4 Representative examples of motion artifacts in en-face OCTA images from two different patients. Vertical stripes partially obscure the underlying microvascular network. 21
2.5 Overview of the OCT processing and feature extraction pipeline. The main workflow includes data reconstruction and preprocessing, volume normalization and realignment, skin-surface and DEJ segmentation, and feature extraction. The branch marked by (*) was applied only to acquisitions affected by partial or complete skin–probe detachment. For these volumes, a NaN exclusion mask was introduced to exclude unreliable regions from surface estimation, followed by glass-layer artifact removal, before proceeding with the main workflow. 25
2.6 Representative example of macroscopic skin-surface inclination in the original OCT volume. (a) Structural OCT B-scan showing the tilted skin surface. (b) Orthogonal cross-sectional view illustrating the corresponding surface curvature. 27

2.7 Preliminary skin-surface tracking in a representative structural OCT B-scan. (a) B-scan after WLS edge-preserving filtering. (b) Preliminary surface detected by dynamic programming and superimposed on the filtered B-scan. (c) Same preliminary surface superimposed on the corresponding original B-scan. 29
2.8 Representative correspondence between an uninformative OCTA region and the associated exclusion mask. (a) Depth-color-coded en-face OCTA projection showing the region affected by signal loss. (b) Corresponding exclusion mask, where violet indicates valid tissue regions and yellow identifies areas affected by tissue detachment. . . 30
2.9 Representative structural OCT B-scan before and after macroscopic surface realignment, showing the correction of large-scale tissue inclination while preserving local skin-surface topography. 32
2.10 Glass artifact suppression. Left: Identification of the high-intensity glass reflection interface and defined masking margins. Right: Cleaned volume with the glass reflection suppressed. 33
2.11 Schematic representation of the developed graph-based skin-surface segmentation pipeline. For each B-scan, edge- and region-based information is combined into a cost map, which is converted into a directed graph and optimized through a max-flow/min-cut approach to identify the skin surface. 34
2.12 Gradient computation for the edge-based cost. (a) Representative preprocessed structural OCT B-scan. (b) Horizontal gradient component g_x obtained using the Scharr operator. (c) Vertical gradient component g_y obtained using the Scharr operator. 35
2.13 Representative edge-based cost map obtained by combining the normalized Scharr gradient magnitude with the orientation-dependent weighting function. Low-cost regions preferentially correspond to strong dark-to-bright transitions compatible with the expected skin-surface boundary. 36
2.14 Representative construction of the region-based cost. (a) Structural OCT intensity map. Pixels exceeding the selected intensity threshold ($\theta = 100$) were identified as bright tissue pixels. (b) Cumulative bright-pixel map obtained by integrating the thresholded binary values along the axial direction for each A-line. The resulting map provides the region-based contribution used for skin-surface segmentation. 37
2.15 Refined skin-surface segmentation. (a) Representative structural OCT B-scan with the segmented skin surface superimposed in red. (b) Three-dimensional representation of the resulting skin-surface map. 38

2.16 Representative DEJ detection from the local axial intensity profile. The first two significant intensity peaks (p_1 and p_2) are identified within the search window, and the DEJ candidate is assigned to the minimum-intensity position between them. 40
2.17 Skin-surface and Dermo-Epidermal Junction segmentation. (a) Representative structural OCT B-scan showing the segmented skin surface and DEJ. (b) Three-dimensional representation of the detected DEJ surface. 41
2.18 Representative effect of the NaN exclusion mask on the segmented skin and Dermo-Epidermal Junction (DEJ) surface maps. (a) Three-dimensional skin-surface map before application of the exclusion mask. (b) Corresponding skin-surface map after removal of the masked regions. (c) Corresponding DEJ surface map after removal of the masked regions. 41
2.19 Skin-surface roughness analysis. (a) Depth map obtained from the segmented skin surface, where pixel values represent the detected axial surface position. (b) Corresponding Frangi-filter response used to identify surface furrows. 43
2.20 Schematic representation of the longitudinal classification analysis workflow. The process includes data preprocessing, feature selection, dimensionality reduction, and multiclass classification using an XGBoost model. 45
2.21 Schematic representation of the toxicity-classification workflow, including patient-grouped cross-validation, stability-based feature selection, and evaluation of Logistic Stable, Ridge, and ElasticNet classifiers. 48

3.1 Representative comparison between manual and automatic vascular segmentation in an OCTA acquisition characterized by good image quality. (a) Original OCTA en-face projection. (b) Manual vascular segmentation used as reference. (c) Corresponding segmentation obtained using the automatic network. 51
3.2 Representative comparison between manual and automatic vascular segmentation in an OCTA acquisition affected by motion artifacts. (a) Original OCTA en-face projection. (b) Manual vascular segmentation used as reference. (c) Corresponding segmentation obtained using the automatic network. 51
3.3 Representative OCTA en-face projection affected by motion artifacts before and after correction. (a) Original acquisition showing prominent stripe-like motion artifacts. (b) Corresponding OCTA projection after application of the motion-artifact correction procedure. 52

3.4 Representative comparison of vascular segmentation for an OCTA acquisition affected by motion artifacts. (a) Manual reference segmentation. (b) Automatic segmentation of the original motion-corrupted volume. (c) Automatic segmentation obtained after motion-artifact correction. 53
3.5 Representative challenging OCTA acquisition showing a marked discrepancy between manual and automatic vascular segmentation. (a) Original en-face OCTA projection. (b) Corresponding depth-color-coded representation. (c) Manual reference segmentation. (d) Automatic segmentation of the original acquisition. (e) Automatic segmentation after application of the motion-artifact correction procedure. 54
3.6 Normalized confusion matrices obtained from the longitudinal classification analysis using features derived from: (a) manual reference segmentations; (b) original automatic segmentations; and (c) optimized automatic segmentations after targeted motion-artifact correction. Each row is normalized by the number of observations belonging to the corresponding true class. 56
3.7 Representative outputs of the OCT morphological processing pipeline for the same acquisition. (a) Original OCT B-scan. (b) Corresponding B-scan with the detected skin surface and dermo-epidermal junction (DEJ) superimposed. (c) Three-dimensional representation of the segmented skin surface. (d) Three-dimensional representation of the segmented DEJ surface. (e) Surface-depth map derived from the reconstructed skin morphology. (f) Frangi-derived map used for the characterization of the superficial furrow pattern. 57
3.8 Representative example of non visible DEJ in a B-scan 58
3.9 Representative OCT B-scans showing showing challenging DEJ visualization. (a) Interruption of the DEJ. (b) Irregular and interrrupted DEJ. (c) Hyporeflective areas below the DEJ. 58
3.10 Representative OCT B-scans showing crusts structure. 59
3.11 Representative outputs of the OCT morphological processing pipeline for the same acquisition. (a) Original OCT cross-scan section. (b) Original B-scan section after glass artifact removal. (c) B-scan with the detected skin surface and dermo-epidermal junction (DEJ) superimposed. (d) Three-dimensional representation of the segmented skin surface after NaN mask application. (e) Three-dimensional representation of the segmented DEJ surface after NaN mask application. (f) Surface-depth map derived from the reconstructed skin morphology. (g) Frangi-derived map used for the characterization of the superficial furrow pattern. 60

3.12 Representative OCT B-scan showing partial detachment area with the glass (a) B-scan before glass artifact removal (b) B-scan after glass artifact removal 61
3.13 Distribution of OCT-derived morphological features at PRE, POST, 3M, 6M, and in healthy skin. 62
3.14 Longitudinal evolution of OCT-derived morphological features according to toxicity class. Healthy skin measurements are reported separately as a reference and are not connected to the longitudinal trajectories. 63
3.15 Individual and group-level longitudinal evolution of OCT-derived morphological features according to toxicity class. Thin lines represent individual lesion trajectories, while thick lines and markers represent the corresponding group-level values for toxicity classes 0 and 1. Healthy skin measurements are displayed separately as an independent reference. 64
3.16 Normalized confusion matrices obtained from the longitudinal classification analysis using: (a) OCTA-derived microvascular and radiomic features (OCTA/Radiomics); and (b) the extended feature set additionally including OCT-derived morphological features (OCT/OCTA/Radiomics). 65
3.17 Ranking of the highest-contributing features in the longitudinal classification model using the extended OCT/OCTA/Radiomics feature set. 66
3.18 Permutation analysis of the PRE + Δ3M ridge toxicity classifier using the extended OCT/OCTA/Radiomics feature set. The distribution represents the ROC-AUC values obtained after random permutation of the toxicity labels, while the vertical line indicates the ROC-AUC obtained using the original labels (ROC-AUC = 0.88; permutation $p = 0.0099$). 68
3.19 Paired bootstrap distributions of the differences in ROC-AUC between the OCT/OCTA/Radiomics and OCTA/Radiomics feature sets. Panels (a–c) correspond to PRE, panels (d–f) to PRE + ΔPOST, and panels (g–i) to PRE + Δ3M. Within each row, Logistic Stable, Ridge, and Elastic Net classifiers are shown from left to right. Positive ΔAUC values indicate higher ROC-AUC after inclusion of OCT-derived morphological features, whereas negative values indicate higher ROC-AUC for the OCTA/Radiomics feature set. . . 70

4.1 Representative OCT B-scans highlighting morphological alterations discussed in the text. Red circles indicate regions in which the dermo-epidermal junction (DEJ) appears locally interrupted or poorly defined, while red arrows indicate hyporeflective dermal areas. . . . 75

Chapter 1

Introduction

1.1 Clinical Background: Non-Melanoma Skin Cancer

Non-Melanoma Skin Cancer (NMSC), also called keratinocyte carcinoma, is the most common type of cancer in the Caucasian population [1].The development of NMSC is caused by a mix of genetic, phenotypic, and environmental factors.

The most important factors that could affect the incidence of non-melanoma skin cancer development are age, gender, and sex. The risk of developing NMSC, in fact, increases with age, has a higher incidence in men than in women, and has a strong correlation with solar irradiance, thus it has a much higher probability of growing on sun-exposed body sites.

1.1.1 Characterization of BCC, SCC, and AK

The term NMSC includes some rare skin tumors (like cutaneous lymphomas or Merkel-cell carcinomas), but it is mainly used to define Basal Cell Carcinoma (BCC) and Squamous Cell Carcinoma (SCC). Another important lesion is Actinic Keratosis (AK), which is a pre-cancerous precursor.

- **Basal Cell Carcinoma (BCC):** It represents about 75% to 80% of all NMSC cases. It starts from basal keratinocytes in the epidermis and is linked to intermittent sun exposure, especially during childhood. BCC grows slowly and has an extremely low metastatic rate ($< 0.1\%$) [2]. Clinically, BCC can look very different: it shows pearly nodules with telangiectasia, ulcerated plaques, or pigmented areas. Histologically, it is divided into subtypes like superficial, nodular, micronodular, infiltrative, and morphoeic. Nodular BCC is the most common (over 60% of cases). Infiltrative and micronodular subtypes are more aggressive and tend to recur.

- **Squamous Cell Carcinoma (SCC):** It accounts for most of the remaining NMSC cases. It may arise from precursor lesions such as actinic keratosis (AK) and Bowen's disease, and represents a multistep accumulation of genetic damage. It is directly linked to chronic, cumulative lifetime sun exposure. Unlike BCC, SCC has a significant metastatic rate between 0.3% and 3.7%. The risk of metastasis is higher for tumors larger than 2 cm, thicker than 2–4 mm, poorly differentiated, or growing on chronic scars and old radiation sites.

1.1.2 Radiotherapy Management and Treatment Alternatives

The choice of treatment depends on the risk of the tumor, patient health, and available services. The primary goal is to completely remove the tumor, while the secondary goal is to get a good aesthetic result. Treatments are divided into surgical and medical options [1].

Among the surgical interventions, the standard surgical excision is the most common treatment for well-defined tumors, followed by Mohs micrographic surgery, which represents the gold standard for high-risk, aggressive or recurrent tumors that appear especially on the face. In case of small, well-differentiated, primary slow-growing tumors usually other destructive techniques are exploited, such as curettage and electrodesiccation. Furthermore, cold-induced destruction of non-melanoma skin cancer, achieved with liquid nitrogen cryosurgery, is considered an appropriate technique for treatment of low-risk primary basal cell and squamous-cell carcinomas.

Medical and Non-Surgical Alternatives

Topical treatments (like 5-fluorouracil cream or 5% imiquimod cream) and photodynamic therapy (PDT) are excellent for superficial BCC and AK because they give great cosmetic results. However, their 5-year cure rate for superficial BCC is lower than surgery (around 78%–84%). They are not recommended for aggressive BCC or invasive SCC because the risk of recurrence is too high.

Radiotherapy (RT)

Superficial radiotherapy, electron beam radiotherapy, and brachytherapy are highly effective for patients who cannot or do not want to undergo surgery. RT shows a 5-year cure rate of about 91.3% for BCC and 90% for SCC. It is also used as an adjuvant treatment after surgery when margins are positive, or in cases of perineural invasion and lymph node metastases. On the lip, ear, and nose, RT can provide the best cosmetic and functional results.

However, RT can damage healthy skin over time, causing side effects like skin atrophy, necrosis, dyspigmentation, telangiectasia, and radiodystrophy. Furthermore, RT increases the risk of developing new skin cancers after 15–20 years. For this reason, it should be used cautiously in young patients (< 65 years), and it is completely contraindicated on areas already treated with radiation or in patients with specific genetic syndromes.

1.1.3 Radiation-Induced Skin Toxicity

Although radiotherapy represents one of the primary treatment modalities for NMSC, its ionizing effects inevitably impact the healthy surrounding tissues, leading to radiation-induced skin reactions (RISR), clinically referred to as radiation-induced skin toxicity or generally adverse events (AE). The clinical manifestations of this toxicity range from acute reactions, such as erythema and edema, to severe chronic alterations, including ulceration and fibrosis. These adverse reactions cause significant patient discomfort, severely impairing their quality of life, and can lead to clinical complications severe enough to cause the temporary or permanent interruption of the oncological treatment.

To systematically classify and grade the severity of these cutaneous AE, standardized clinical grading scales are used. In this study, the Common Terminology Criteria for Adverse Events (CTCAE), established by the National Cancer Institute (NCI), was adopted as the reference framework [3]. Within this system, an AE is defined as any unfavorable and unintended sign, symptom, or disease temporally associated with the use of a medical treatment or procedure, regardless of whether it is considered causally related to the treatment itself.

A severity scale is provided for each AE term, varying from 1 through 5 with an increasing level of clinical severity:

- **Grade 1 (Mild):** Asymptomatic or mild symptoms; clinical or diagnostic observations only; intervention is not indicated.
- **Grade 2 (Moderate):** Minimal, local, or non-invasive intervention indicated; limiting age-appropriate instrumental ADL[1].
- **Grade 3 (Severe or medically significant):** Not immediately life-threatening; hospitalization or prolongation of hospitalization indicated; disabling; limiting self-care ADL[2].

[1]Instrumental Activities of Daily Living (iADL) refer to preparing meals, shopping for groceries or clothes, using the telephone, managing money, etc.

[2]Self-Care Activities of Daily Living (sADL) refer to bathing, dressing and undressing, feeding oneself, using the toilet, taking medications, and not being bedridden.

- **Grade 4 (Life-threatening consequences):** Urgent intervention indicated to prevent death.
- **Grade 5 (Death):** Death directly related to the adverse event.

Within the scope of this study, the CTCAE scale is uniquely suited for assessing *Skin and subcutaneous tissue disorders*, with particular attention paid to skin hyperpigmentation/hypopigmentation, erythema, radiation dermatitis (encompassing dry and moist desquamation), and telangiectasia. Consequently, this thesis investigates how these microvascular and morphological parameters of the skin tissue correlate with the specific grade of skin toxicity developed by each patient during or after the radiotherapy protocol.

1.2 Optical Coherence Tomography (OCT)

Optical Coherence Tomography (OCT) is a completely non-invasive diagnostic imaging modality designed to generate high-resolution, cross sectional depth resolved 2D and 3D images of biological tissue *in vivo*. Its foundational operating principle consists of analyzing the time delay and the backscattering intensity of the waves reflected by a tissue specimen to differentiate structural interfaces.

1.2.1 Physical Principles and Signal Acquisition

Since the velocity of light precludes the direct electronic measurement of temporal echo delays, an optical reference framework must be established. This framework is implemented via low-coherence interferometry, typically executed through a Michelson-type or a Mark-Zehnder interferometer architecture. A standard bulk-optics or fiber-based Michelson interferometer consists of a light source, a beam splitter (or fiber coupler), a reference arm, a sample arm, and a photo-detection system (Figure 1.1).

The beam splitter divides the emission from a broadband, low-temporal-coherence light source into two distinct optical paths:

- **In the sample arm:** The probing beam is focused into the specimen, where it undergoes back-scattering and reflection at various depth layers.
- **In the reference arm:** The beam is directed toward a reference mirror positioned at an optical path length, which is close to the sample arm path length.

Upon reflection from their respective paths, the two beams return to the beam splitter/coupler to physically recombine at the interferometer output before striking the photodetector. The physical superposition of these electromagnetic field

amplitudes generates optical interference. Coherent waves add constructively or destructively depending on their relative phase relationship.

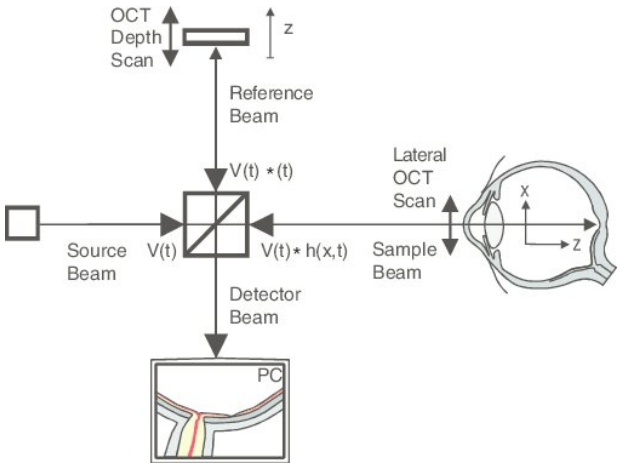


Figure 1.1: Standard OCT system based on Michelson architecture. LS=low time-coherence light source, PC= personal computer [4]

Due to the low temporal coherence of the broadband source, a fixed phase relation is maintained only over a highly restricted travel range called the coherence length (l_c), which acts as a microscopic depth-selection gate. Assuming a source with a Gaussian-shaped power spectrum, the round-trip coherence length l_c in air is mathematically defined as:

$$\delta z = l_c = \frac{2\ln 2}{\pi} \cdot \frac{\lambda_0^2}{\Delta\lambda_{\mathrm{FWHM}}} \tag{1.1}$$

where λ_0 represents the central emission wavelength and $\Delta\lambda_{\mathrm{FWHM}}$ denotes the spectral bandwidth at full width at half maximum.

If a narrow spectral bandwidth (high coherence) light source is used, the interference signal remains active across a wide spatial range, making it impossible to isolate individual depth layers.

Instead, a broadband light source exhibits low temporal coherence. As a result, constructive interference can occur exclusively within a highly confined micrometer-scale volume, precisely where the optical path lengths of the reference and sample arms are perfectly matched. This strict spatial gating is what enables high-resolution depth localization (A-scan profiling). Guiding the probing beam laterally across the tissue compiles a sequential series of A-scans into a cross-sectional brightness profile known as a B-scan.

1.2.2 OCT Techniques and Parameters

OCT technology has undergone significant architectural transitions, primarily categorized into two distinct generations: Time Domain OCT (TD-OCT) and Fourier Domain OCT (FD-OCT).

In first-generation TD-OCT configurations, depth-resolved information is captured sequentially in time. To acquire a single A-scan, the reference mirror must be mechanically translated along a longitudinal stage to scan the coherence gate through the depth of the sample. The back-reflected light from each arm is combined, but it can only interfere to produce an intensity modulation if the path lengths match. This mechanical translation severely limits data acquisition speeds and results in a lower signal-to-noise ratio (SNR) that is highly susceptible to motion artifacts.

Mathematically, the total time-averaged intensity $I_E(t;\Delta t)$ striking the photodetector as a function of the echo time delay Δt can be modeled as [4]:

$$I_E(t;\Delta t) = I_S(t) + I_R(t) + G_{SR}(\Delta t) \tag{1.2}$$

Here, $I_R(t)$ is the constant DC intensity returning from the reference mirror, $I_S(t)$ is the constant DC intensity returning from the scattering structures within the sample arm, and $G_{SR}(\Delta t)$ is the real-valued cross-correlation or interference term carrying the structural data. The cross-correlation term represents the physical interference between the sample and reference fields and is expressed as [4]:

$$G_{SR}(\Delta t) = 2\mathrm{Re}\left\{V_S^*(t)V_R(t+\Delta t)\right\} = 2\mathrm{Re}\left\{\Gamma_{SR}(\Delta t)\right\} \tag{1.3}$$

where $\Delta t = \frac{2(z_R - z_n)}{c}$ represents the temporal delay resulting from the path difference between the reference mirror position (z_R) and the n-th sample layer (z_n). Utilizing the analytic signal representation, the complex cross-correlation function $\Gamma_{SR}(\Delta t)$ can be related to its real parts via the Hilbert Transform (HT) [4]:

$$\Gamma_{SR}(\Delta t) = \frac{1}{2}G(\Delta t) + \frac{i}{2}\mathrm{HT}\left\{G(\Delta t)\right\} \tag{1.4}$$

Second-generation FD-OCT overcomes the mechanical constraints of TD-OCT by shifting data acquisition to the spectral domain. By measuring the spectral components of the interferometric signal over a stationary reference arm, FD-OCT captures information from all depth layers simultaneously, providing a dramatic increase in both acquisition speed and SNR.

The total power spectrum $S(\nu;\Delta t)$ detected in the frequency domain is modulated by rapid periodic oscillations and is expressed as [4]:

$$S(\nu;\Delta t) = S_S(\nu) + S_R(\nu) + 2\mathrm{Re}\left\{W_{SR}(\nu)\right\}\cos(2\pi\nu\Delta t) \tag{1.5}$$

where $S_S(\nu)$ and $S_R(\nu)$ are the individual power spectra, and $W_{SR}(\nu)$ is the cross-spectral density, which represents the direct Fourier Transform (FT) of the complex mutual coherence function $\Gamma_{SR}(\Delta t)$ [4]:

$$W_{SR}(\nu) = \mathrm{FT}\left\{\Gamma_{SR}(\Delta t)\right\} \tag{1.6}$$

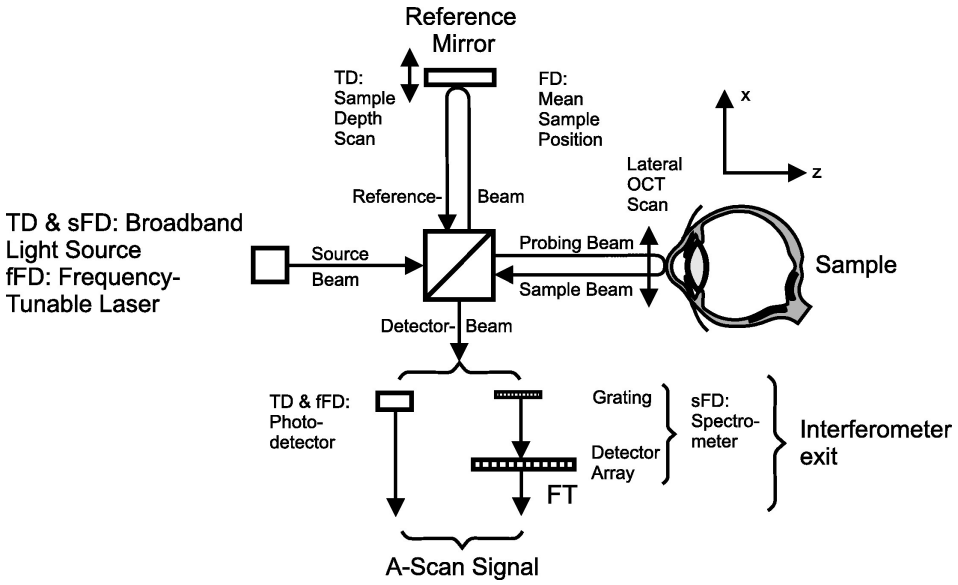


Figure 1.2: Standard time-domain (TD) and Fourier-domain (FD) Michelson interferometer configurations used in OCT. [5]

FD-OCT can be further divided into two distinct implementations based on the light source and detection strategy employed [6]. Both of them are represented in Figure 1.2 and are described below:

- **Spectral Domain OCT (SD-OCT):** Utilizes a static broadband light source (such as a superluminescent diode, SLD) and a high-speed spectrometer with a diffractive element that spatially separates the different wavelength contributions into a line image. This image is then instantaneously recorded by a high speed line scan camera.
- **Swept-Source OCT (SS-OCT):** Employs a specialized laser source that sweeps a narrow line-width over a broad range of wavelengths. For each sweep, a wavelength component of the interferometric signal is detected sequentially by a high speed photodetector. SS-OCT is capable of mitigating the signal sensitivity drop with depth (known as sensitivity roll-off) that typically limits SD-OCT systems.

The main parameters of an OCT instrument are its optical power and wavelength, its penetration, resolution and sensitivity. Whereas these parameters are interdependent, some might be determined by the imaging problem.

Sensitivity and Depth Limits: Sensitivity can be defined as the ratio between the reflectivity of a perfectly reflecting mirror ($R = 1$) and the smallest sample reflection coefficient (R_{min}) that yields a signal power equal to the noise floor, corresponding to a signal-to-noise ratio (SNR) of 1 [5]. Mathematically, it is expressed as:

$$S = \frac{1}{R_{\text{min}}} \tag{1.7}$$

In decibels (dB), this is commonly represented as:

$$S_{[\text{dB}]} = 10 \log_{10} \left(\frac{1}{R_{\text{min}}} \right) \tag{1.8}$$

The maximum available imaging depth (z_{max}) in frequency-domain systems is fundamentally capped by the spectral resolution or finite sampling capacity of the detection apparatus. It is defined according to the Nyquist-Shannon criteria as [6]:

$$z_{\text{max}} = \frac{1}{4} \cdot \frac{\lambda_0^2}{\frac{\Delta\lambda}{N}} \tag{1.9}$$

where N represents the number of recorded spectral sample points across the full bandwidth width $\Delta\lambda$.

Unlike the imaging depth limit, the real *penetration depth* is a tissue-dependent parameter. For highly scattering dermatological tissues, a central wavelength of 1300 nm is preferred within the near-infrared spectrum to optimize the compromise between deep tissue penetration and viable axial resolution. By using this wavelength, the penetration depth reaches values of $1 - 1.5$ mm.

Axial Resolution (δz): The axial resolution is exclusively governed by the coherence properties of the light source and is entirely independent of the system's focusing optics. Corrected for a biological medium with a refractive index n, the axial resolution inside the tissue is expressed as [6]:

$$\delta z_{\text{tissue}} = \frac{2 \ln 2}{\pi} \cdot \frac{\lambda_0^2}{n \cdot \Delta\lambda_{\text{FWHM}}} \tag{1.10}$$

Lateral Resolution (δx): Instead, the lateral resolution depends strictly on diffraction-limited focusing optics and the beam geometry entering the objective.

Assuming a Gaussian beam profile, the lateral spot size at half maximum intensity is given by [6]:

$$\delta x = \frac{\sqrt{2\ln 2} \cdot 2 f_{\text{sys}} \lambda_0}{\pi d} = \frac{\sqrt{2\ln 2} \cdot \lambda_0}{\pi \cdot \text{NA}} \tag{1.11}$$

where f_{sys} is the focal length of the optical system, d is the entering beam diameter at the $1/e^2$ intensity decay point, and NA is the numerical aperture. Tight focusing would result in a higher lateral resolution, but at the same time it reduces the depth of focus, so the axial range over which the laser beam remains tightly focused, which is obtaines by [6]:

$$b = \frac{n \cdot \lambda_0}{2\pi \cdot \text{NA}^2} \tag{1.12}$$

1.2.3 Morphological OCT vs. Functional OCT-Angiography (OCTA)

While standard structural OCT captures the static, morphological cross-sectional aspect of the scanned tissue, Optical Coherence Tomography Angiography (OCTA) uses structural intensity variance as a functional foundation to isolate microvascular perfusion without requiring exogenous contrast dyes. This avoids invasive, time-consuming procedures and potential side effects associated with external contrast agents.

To enhance the microvasculature over the static background tissue, the OCTA technique rapidly performs repeated B-scans at the exact same tissue location over a designated short time interval. Motion contrast is then obtained by calculating the decorrelation or dissimilarity between pairs of consecutive images.

The continuous motion of scattering red blood cells within vessels induces significant temporal fluctuations in both the intensity and phase components of the complex-valued OCT signal. Consequently, to quantify this dissimilarity, the variation in intensity or phase is calculated, and the mean of these dissimilarities is computed to generate high-resolution maps of the microvascular network.

1.2.4 State of the Art in OCT/OCTA Analysis

OCT and OCTA have been investigated in different clinical applications, including ophthalmology, gastroenterology, and dermatology [4]. Ophthalmology represents one of the earliest and most established fields of application, where these techniques have been used for the visualization of retinal structures and microvascular networks [7]. Gastroenterology has also benefited from these techniques, particularly in the assessment of gastrointestinal mucosa and other structures and early detection of neoplasms through high-resolution images [4].

In dermatology, OCT and OCTA have been investigated for the characterization of skin lesions and disorders, mapping of tissue structures, detection and monitoring of skin tumors, and assessment of inflammatory skin conditions [4, 8, 9]. In addition to qualitative image assessment, increasing attention has been directed toward the extraction of quantitative parameters capable of objectively describing tissue structure and microvascular organization.

From a structural perspective, epidermal thickness represents one of the most investigated parameters. Its physiological variability has been characterized across different anatomical regions and subject characteristics, while changes in epidermal morphology have also been associated with skin alterations and pathological conditions [10, 11, 12, 13].

Skin-surface morphology and roughness have also been investigated as quantitative descriptors, with potential applications in the characterization of skin aging and treatment-related changes [14]. OCT is particularly attractive in this context because it enables non-invasive, in vivo assessment of epidermal morphology.

OCTA provides instead functional visualization of the microvascular network, enabling the extraction of quantitative parameters such as vessel density, length, diameter, organization and tortuosity, which have been investigated as potential descriptors of pathological tissue [15, 16].

Studies demonstrated characteristic vascular alterations measured through OCTA in different skin conditions [17, 8]. In inflammatory lesions, vascular dilation and increased local perfusion were observed, whereas basal cell carcinomas exhibited denser and more disorganized vascular networks, with relatively large vessels located close to the skin surface [17].

These observations illustrate how structural and vascular imaging can provide complementary information on the same tissue.

The extraction of these quantitative descriptors strongly depends on the accurate identification of the anatomical structures and vascular networks of interest, making segmentation a fundamental step of the image-analysis pipeline. Manual segmentation has traditionally been used as a reference approach [9]; however, it is time-consuming, operator-dependent, and difficult to apply to large three-dimensional datasets. These limitations have motivated the development of semi-automated and fully automated approaches.

For structural OCT, different segmentation strategies have been proposed. Conventional approaches include shapelet-based methods for epidermal boundary detection [18], graph-based methods for skin-surface segmentation [14], attenuation-based graph optimization for DEJ localization [19], and local integral projection methods for epidermal boundary identification [20]. Image flattening has also been employed after skin-surface detection to compensate for tissue curvature and restrict the search region for deeper interfaces, thereby facilitating subsequent boundary localization [19].

More recently, learning-based approaches have been introduced to further automate OCT segmentation. Convolutional neural networks (CNNs) can learn relevant image representations from annotated datasets and can also be combined with conventional optimization methods. For example, a CNN–graph-search framework has been proposed for rapid epidermal boundary identification and thickness measurement [21]. While these approaches can provide accurate segmentation, computational requirements and the availability of suitable annotated datasets remain important considerations for their practical implementation.

Automatic segmentation is equally important for OCTA, where accurate vascular delineation is required for quantitative analysis. OCTA images can be affected by motion and acquisition-related artifacts which could limit the reliability of simple thresholding or morphological approaches. Consequently, accurate pipelines have been developed to mitigate these artifacts and improve the reliability of vascular segmentation [9, 22]. Learning-based approaches, including two and three-dimensional U-Net architectures, have also been investigated to automate vascular segmentation [9].

Overall, previous studies demonstrate the potential of quantitative OCT and OCTA analysis for extracting complementary structural and vascular information from skin. Nevertheless, robust multimodal analysis remains challenging, especially in heterogeneous clinical datasets.

1.3 Purpose of the Study

This study is embedded within a broader clinical research framework aimed at evaluating microvascular alterations in patients affected by Non-Melanoma Skin Cancer undergoing a 6-month radiotherapy protocol. Previous longitudinal assessments using OCTA successfully demonstrated the technique's capability to quantitatively monitor vascular remodeling and differentiate between distinct treatment timepoints only based on vascular pattern analysis.

Despite these promising preliminary results, significant methodological limitations remained, including:

- **Motion Artifacts:** High sensitivity to patient movement, which can degrade image quality and obscure microstructural details.
- **Manual Segmentation:** A highly time-consuming, operator-dependent process requiring dedicated software, which limits reproducibility and clinical throughput.

To overcome these challenges and enhance the robustness of the existing dataset, this thesis introduces two technical improvements:

- **Advanced Artifact Mitigation:** The application of a pre-developed motion artifact removal algorithm to improve image fidelity.
- **Automated Segmentation:** The use of a pre-developed automated pipeline to segment the vasculature, reducing processing time and minimizing inter-operator variability.

Furthermore, to provide a more complete evaluation, the angiographic data has been complemented by a structural OCT analysis. This dual approach focuses on the morphological characterization of the skin, spanning from superficial epidermal layers to the deeper dermis.

Ultimately, this study aims to combine structural and functional (vascular) parameters into a comprehensive, multi-modal analytical framework. Validating this integrated approach represents a crucial step toward providing dermatologists with clear, quantifiable, and clinically interpretable biomarkers. This tool will support clinicians in monitoring lesion regression, assessing real-time tissue status, and predicting potential long-term radiation-induced toxicities.

Chapter 2

Materials and Methods

2.1 Dataset and Clinical Protocol

2.1.1 Patient Cohort and pre-existing Framework

The clinical dataset utilized in this thesis was originally collected during an Institutional Review Board (IRB)-approved longitudinal study conducted at the Center for Medical Physics and Biomedical Engineering of the Vienna General Hospital (Medical University of Vienna, number 1246/2013). A total of 20 patients diagnosed with Non-Melanoma Skin Cancer, specifically BCC, SCC, or AK, were enrolled in the trial. The diagnosis was histologically confirmed through biopsy for all lesions prior to the therapeutic phase. Following a multidisciplinary evaluation, patients for whom surgical intervention was contraindicated were referred for superficial orthovoltage radiotherapy. To evaluate the microvascular and morphological alterations induced by ionizing radiation, patients underwent imaging sessions at four specific timepoints, integrated into their routine clinical follow-up, over a 6-month follow-up period: prior to the start of radiotherapy (Pre), immediately following the final radiation fraction (Post), 3 months post-treatment (3M), and 6 months post-treatment (6M). A baseline scan (Healthy) was also acquired on a non-cancerous area, close to the lesion, to approximate the healthy local vascular structure.

The initial dataset, is composed of 178 volumes, divided in the different timepoints as shown in Table 2.1. One patient was excluded from the analysis due to technical issues with the imaging data.

Table 2.1: Overview of the longitudinal clinical dataset distribution across treatment timepoints.

Cohort Size	Healthy Baseline	Pre-RT	Post-RT	3M	6M
19 Patients	26 Volumes	43 Volumes	40 Volumes	40 Volumes	29 Volumes

2.1.2 Previous processing state

The initial database available for this study comprised raw interferometric volumes which had undergone a preliminary preprocessing pipeline developed in previous works. This foundational pipeline included [15]:

- 3D reconstruction of the structural OCT volumes from raw B-scans.
- Extraction of the functional OCTA volumes via an intensity-based decorrelation algorithm, designed to isolate the patient microvasculature by filtering out the surrounding static tissue matrix.
- Manual 3D segmentation using AMIRA (2024, Thermo Fisher Scientific Inc., Waltham, MA, USA) to yield binary vascular masks, where voxels representing perfusion are assigned a value of 1, and the background is set to 0.

During these preliminary clinical assessments, specific volumes exhibited regional avascular areas, superficial blisters, or optical shielding caused by highly exophytic squamous crusts. Since these regions lacked viable angiographic information and could severely bias the subsequent quantitative feature extraction, they were reviewed alongside trained radiotherapists. Consequently, dedicated binary Not-a-Number (NaN) masks were manually generated in AMIRA to assign a NaN value to these uninterpretable regions, effectively excluding them from mathematical analysis without affecting the surrounding valid tissue parameters. These masks and segmentations were stored in a structured subfolder architecture within the main repository.

Following these steps, a total of 50 OCTA features and 107 radiomic features were extracted to obtain a quantitative representation of the analyzed volumes. While the OCTA features characterized vascular density, tortuosity, and morphological organization of the segmented vasculature, the radiomic features described the intensity distribution, texture, and spatial organization of the OCT signal. Together, these features were merged and used to train an XGBoost classifier to assess their ability to distinguish among the different timepoints [16].

2.1.3 Radiotherapy Protocol

The radiation was delivered using a Gulmay D3300 superficial orthovoltage unit (Gulmay Medical Ltd., Chertsey, UK) [16]. This system works with a bi-polar ceramic X-ray tube, featuring an 8 mm focal spot and a tungsten target. Its built-in filtration consists of 3 mm of beryllium, which allows the lowest-energy X-rays to pass through, making it ideal for treating superficial lesions.

To adapt the treatment to the specific depth and size of each tumor, the tube potential and current can be adjusted flexibly between 80–250 kVp and 0–30 mA, respectively. This flexibility allowed for the customization of the beam energy for each patient, ensuring the radiation reached the correct depth without unnecessarily exposing deeper healthy tissues. The tumor was targeted with a safety margin ranging from 5 mm to 20 mm, while the surrounding area was protected using a custom-made lead shield [16].

Different dose fractionation schedules were selected based on the patient's age as well as the location, size, and thickness of the tumor. The prescribed radiation protocols are summarized in Table 2.2. One patient did not complete the treatment.

Table 2.2: Overview of the alternative radiation dose fractionation schedules and patient distribution.

Number of Patients	Total Dose (Gy)	Number of Fractions	Dose per Fraction (Gy)
4	50	20	2.5
1	48	16	3.0
1	45	15	3.0
9	40	10	4.0
3	36	6	6.0

2.1.4 Laboratory System

Imaging was performed using a custom-built, high-resolution laboratory SS-OCT system utilizing a specialized laser source from Insight Photonic Solutions, USA. The main technical specifications of the laboratory device are summarized in Table 2.3 [23]:

Table 2.3: Technical specifications and operating parameters of the custom SS-OCT system.

Parameter	Value
Maximum A-scan Repetition Rate	222 kHz
Central Wavelength (λ_0)	1300 nm
Spectral Bandwidth ($\Delta\lambda$)	30 nm
Axial Resolution (in air)	27.3 μm
Lateral Resolution (in air)	31.5 μm
System Sensitivity (SNR)	105 dB

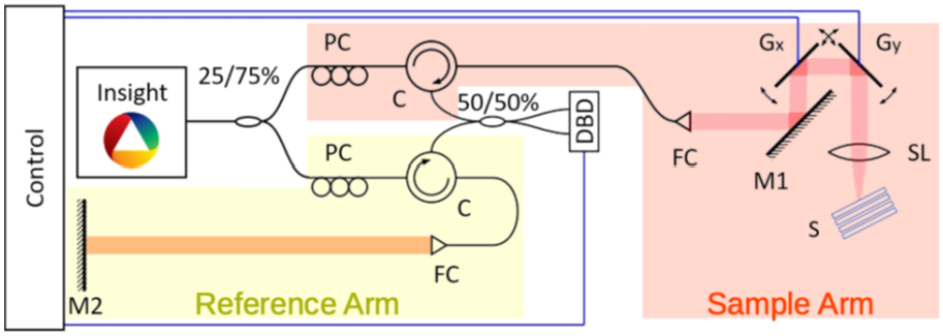


Figure 2.1: Schematic representation of the laboratory SS-OCT system used for clinical acquisitions. Taken from [23]

The system layout is shown in Figure 2.1 [23] . The optical power is split by a fiber coupler, routing $< 75\%$ of the power to the sample arm and $< 25\%$ to the reference arm. The sample arm terminates in a rotatable handheld imaging probe equipped with a flat focusing lens, allowing the operator to adapt the scanning window to different anatomical sites. The reference arm features a translation-stage movable mirror to match the optical path length of the sample arm. Optical circulators (C) are integrated into both paths. These allow guiding the back-reflected light from the arms toward a final 50/50% fiber coupler for recombination, maximizing power efficiency. To suppress laser source common-mode noise and maximize the overall signal-to-noise ratio, a Dual-Balanced Detector (DBD) is utilized at the interferometer exit. After recombination, the light is split equally into two separate fibers entering the two optical inputs of the DBD. The detector electronically subtracts one photodiode signal from the other. This subtraction cancels out the common-mode intensity noise inherent to the laser source while preserving and centering the out-of-phase interferometric cross-correlation term around zero. The resulting analog signal is then routed to an ATS9360 (Alazar

Technologies Inc., Canada) high-speed data acquisition card for digitization. The final system is mounted on a wheeled cart, allowing for easy maneuverability and positioning on the lesion of the patient during clinical acquisitions (Figure 2.2).

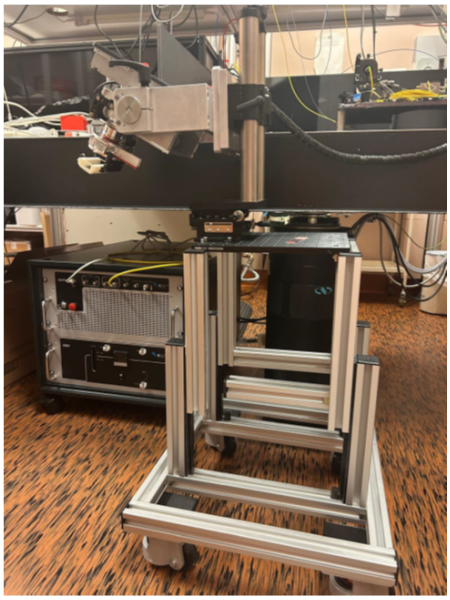


Figure 2.2: SS-OCT system used for clinical acquisitions. Taken from [15]

2.1.5 Scanning Protocol

The clinical acquisition workflow follows a standardized three-phase protocol [15], divided into a series of tasks performed before, during, and after the session:

- **Pre-Acquisition Phase:** This phase consists of laser source calibration to generate a Data Valid Vector (DVV) file, which contains the reliable indices for spectral resampling. This is followed by the sterilization of the patient interface components and the software initialization of a custom LabView acquisition graphical user interface.
- **Intra-Acquisition Phase:** The patient is appropriately positioned on a patient bed, and wears laser safety goggles. The wheeled cart supporting the probe, showed in Figure 2.2 is moved close to the patient, and levers are adjusted for fine-tuning. To achieve spatial stabilization and optimize optical focusing, a standard glass microscope slide mounted on a 3D-printed biocompatible support is attached to the probe tip and placed in direct contact with the skin. Distilled water is introduced at the interface to eliminate air pockets and match the refractive index of the tissue and the imaging window, minimizing superficial specular reflections. The operator utilizes a live B-scan

preview in LabView to ensure optimal alignment before launching the sequence, capturing 4 sequential volumes per position in approximately 20 seconds. The entire acquisition takes from 1 to 2 minutes, due to the time the acquired data needs for saving. In extended lesions more than one acquisition is required, to cover the entire lesion.

- **Post-Acquisition Phase:** This phase involves disposing of the single-use glass slide, disinfecting the system components, saving the digital volumes together with their corresponding calibration files, photographing the lesion, manually labeling the imaged areas of field of views (FoVs) for future refence and recording additional patient clinical notes.

2.1.6 Raw Data Acquisition and Reconstruction

Each acquisition covered an area of approximately $1\,\mathrm{cm}^2$, providing a useful imaging depth of approximately $1-1.5\,\mathrm{mm}$. Although the original acquisition encompassed a larger imaging depth, only the depth range exhibiting adequate image quality was retained during reconstruction, resulting in a final volume of $120 \times 512 \times 512$ voxels, corresponding to approximately $1.15 \times 10 \times 10\,\mathrm{mm}^3$, assuming a skin refractive index of 1.43.

The raw data generated during each acquisition were stored in 32 binary files (B_scans0XX), each containing 64 B-scans. Every B-scan consisted of 512 A-lines, while each A-line contained 2176 depth samples acquired from the detected interference spectrum. However, not all spectral samples were considered reliable. During the calibration procedure, a DVV file was generated to identify the valid spectral region.

OCTA imaging was performed using a BM-scan acquisition protocol, in which four consecutive B-scans were acquired at each lateral position before moving to the next scanning location. These repeated acquisitions enabled the extraction of angiographic information through temporal decorrelation analysis, which exploits the intensity fluctuations generated by moving erythrocytes while static tissue remains highly correlated. Each complete acquisition therefore consisted of 512x512 lateral positions, each sampled four consecutive times.

The temporal interval between two consecutive B-scans at the same lateral position corresponds to the acquisition time of one complete lateral scan added to the galvanometric mirror flyback time. The flyback time represents the interval required by the fast-axis galvanometric scanner to decelerate, reverse its direction, and return to its initial position before the acquisition of the next B-scan.

Starting from these raw spectral measurements, the OCT volumes were reconstructed using an existing Python-based processing pipeline. The reconstruction consisted of reorganizing the binary data into A-lines, retaining only the valid

spectral samples defined by the DVV, removing the background signal by subtracting the lateral mean, applying zero-padding and Gaussian windowing in the spectral domain, and finally performing an inverse Fourier transform to obtain the complex-valued OCT volume used for both structural OCT and OCTA image generation.

2.1.7 Acquisition Challenges and Limitations

A major challenge during the intra-acquisition phase comes from the geometric constraints imposed by the rigid, rectangular profile of the glass slide interface. While uniform mechanical contact and effective skin flattening are easily achievable on planar anatomical regions, such as the extremities or the trunk, maintaining this condition becomes highly problematic on complex topologies. Specifically, when imaging highly curved, narrow, or irregular facial structures (e.g., the nasal bridge, the periorbital region, or the ear), a perfect contact between the glass slide and the cutaneous surface cannot be guaranteed. In these scenarios, the acquisition must be executed either in a contactless mode or with only partial and localized surface contact. This constraint introduces critical technical trade-offs that directly affect both image quality and data consistency:

- **Morphological Inconsistency:** Since the imaged skin is not structurally stretched or stabilized by the slide, the local tissue parameters and the overall morphological appearance within the volume are not perfectly comparable to baseline volumes acquired under full contact conditions.
- **Glass Artifact and Segmentation Decoupling:** In partial contact configurations, the glass slide often remains partially visible within the field of view. This proximity makes it significantly more challenging for automated algorithms to differentiate the superficial glass reflections from the actual skin boundary.
- **Motion Degradation:** The lack of physical stabilization could introduce breathing and micro-motion artifacts that manifest as stripes or blurring in the angiographic volume.

The direct impact of these optical and geometric limitations on the reconstructed 3D volumes will be further analyzed in the following sections.

2.2 Data Optimization Pipeline

To extract quantitative clinical metrics from 3D OCTA volumes, isolating vascular networks through segmentation is a fundamental requirement. In this study,

manual segmentation slice-by-slice through AMIRA software has been considered the clinical reference standard. However, unlike reference annotations obtained under more controlled conditions, this manual segmentation should not be regarded as an error-free ground truth, since the delineation of complex vascular structures is inherently operator-dependent and may be affected by inter- and intra-operator variability. Moreover, manual segmentation is time-consuming when applied to large longitudinal patient datasets. To address the practical limitations of manual segmentation, an advanced GUI-integrated deep learning segmentation network developed at the InVisioning Lab (Biomedical Engineering Group, Politecnico di Torino) was employed. Furthermore, the OCTA volumes, when acquired in vivo, could be affected by several typologies of artifact. Among them, motion artifacts represent one of the most critical issues, leading to severe image degradation and potentially uninterpretable results. To rigorously assess the capabilities of the automated segmentation framework and test the improvements introduced by a motion artifact correction algorithm, a two-step study protocol was designed, shown in Figure 2.3:

1. **Baseline Automated Segmentation:** In the first phase, the segmentation network was applied directly to the raw clinical dataset, perfectly equal to the one used to obtain the manual segmentations. The objective was to test the network's standalone performance and benchmark its boundary predictions against the manual ground truth annotations.

2. **Targeted Artifact Correction and Re-segmentation:** In the second phase, a dedicated preprocessing protocol was introduced to salvage some selected corrupted volumes, which previously had been considered of difficult interpretation. Then, the automatic segmentation was performed again, to obtain the new final features.

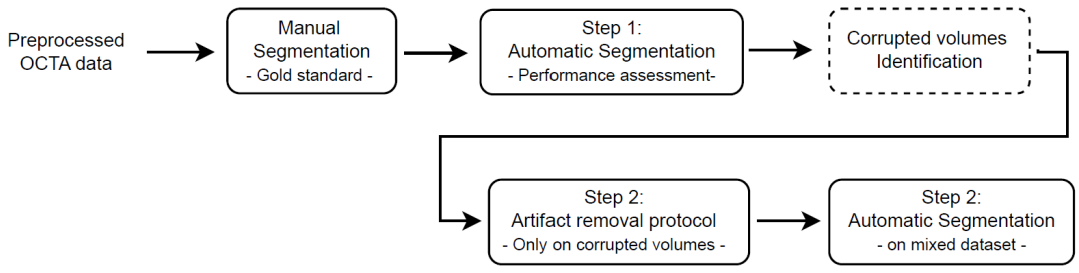


Figure 2.3: Overview of the data-optimization workflow. The analysis comprised two main stages: evaluation of automated vascular segmentation against the manual reference and assessment of the segmentation after motion-artifact correction of selected corrupted OCTA volumes.

Visually, the motion artifacts manifest as vertical stripes overlaid on the en-face projections, which mask the underlying microvasculature. An example is shown in Figure 2.4. Physiologically, this phenomenon is primarily triggered by involuntary patient micro-movements or breathing or could be related to the absence of stable contact between the optical glass and the skin. The algorithm used to optimize these volumes and restore the continuity of the vascular network is based on a Multiresolution Wavelet and Fast Fourier Transform Approach (MW-FFTA) [24], which were applied, as detailed by [22]. The algorithm processes the angiographic volume slice-by-slice through a multi-step mathematical framework:

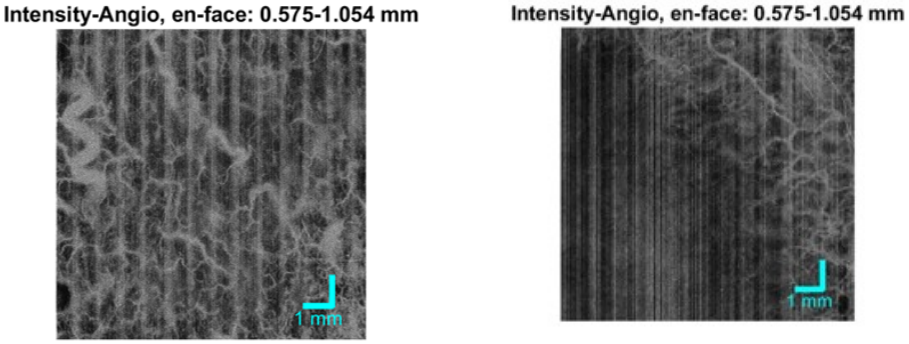


Figure 2.4: Representative examples of motion artifacts in en-face OCTA images from two different patients. Vertical stripes partially obscure the underlying microvascular network.

- **Wavelet Decomposition:** Each en-face slice is decomposed using a 2D Discrete Wavelet Transform (DWT) to isolate directional features, concentrating the vertical striping noise into the vertical detail subbands.
- **Frequency Domain Filtering:** Since the striping noise is highly directional, transforming these vertical subbands into the frequency domain via a 2D Fast Fourier Transform (FFT) concentrates the artifact's energy along a specific horizontal frequency axis. These are selectively attenuated using a customized 2D inverted Gaussian notch mask.
- **Reconstruction and Minimization Constraint:** After returning to the spatial domain via the Inverse FFT (IFFT), the slice is reconstructed using the Inverse DWT. To guarantee that this filtering process does not introduce artificial structures or blurring, a pixel-wise minimization constraint is enforced between the original image and the filtered output. This step guarantees that the algorithm only attenuates the hyper-reflective motion bands without altering the valid vascular morphology.

In order to process only the volumes affected by motion artifact, a subset composed of only those volumes which were presenting the vertical striping noise was created. A selective execution flag (`Motion_Artifact`) was implemented within the execution protocol:

- **`Motion_Artifact = True`:** Assigned to volumes manually identified as corrupted ($N = 101$). Triggers active MW-FFTA stripe attenuation to restore the underlying vascular structure.
- **`Motion_Artifact = False`:** Assigned to artifact-free volumes ($N = 75$). Bypasses filtering entirely, guaranteeing that no unnecessary smoothing or mathematical bias is introduced into clean scans.

The specifics about the volumes affected by motion artifacts are summarized in Table 2.4. The table lists the patient identifiers, lesion indices, and acquisition timepoints for each volume, providing a clear overview of the dataset's quality and the targeted optimization efforts.

Table 2.4: Summary of OCT/OCTA volumes categorized by the presence or absence of motion artifacts across different acquisition timepoints.

Volumes	With Motion Artifact	Without Motion Artifact
ST001.1, ST001.2	post, 6M, healthy	pre, 3M
ST001.3	post, 6M	pre, 3M
ST001.4	post, 3M, 6M	pre, healthy
ST002.1, ST010.1	pre, post, 3M, healthy	—
ST003.1, ST014.1, ST016.1	pre, post, 3M, 6M, healthy	—
ST003.2, ST004.2, ST011.2, ST014.2	pre, post, 3M, 6M	—
ST004.1	post, 3M, 6M	healthy
ST005.1	post	pre, 3M, 6M, healthy
ST005.2	3M	pre, post, 6M, healthy
ST005.3	—	pre, post, 3M, 6M, healthy
ST007.1	pre, post, 3M, 6M	healthy
ST008.1, ST008.2	pre, 3M, 6M, healthy	post
ST009.1	pre, 3M, healthy	post, 6M
ST009.2	pre, healthy	post, 3M, 6M
ST009.3	post	pre, 3M, 6M
ST009.4, ST019.3, ST019.4	—	pre, post, 3M, 6M
ST011.1	pre, 3M, 6M	post, healthy
ST012.1, ST012.2, ST017.3	pre	post
ST012.3	healthy	pre, post
ST013.1	3M, 6M	pre, post, healthy
ST013.2	3M, 6M, healthy	pre, post
ST013.3, ST013.4	6M	pre, post, 3M
ST015.1	pre, post, healthy	—
ST017.1	pre, healthy	post
ST017.2, ST020.1	3M, healthy	pre, post
ST018.1	pre, post, 6M, healthy	3M
ST018.2	pre, 6M	post, 3M
ST019.1	pre, post, healthy	3M
ST019.2	post, 3M	pre, healthy
ST019.5	pre	post, 3M
ST020.2	pre, healthy	post, 3M

The complete dataset, composed of 176 volume files, was saved in NIfTI format (`.nii`), adhering to the standardization requirements of the GUI framework. Files were named as:

`contrast_ST0XXYT.nii`

where `XX` represents the patient identifier, `Y` denotes the lesion index, and `T` indicates the acquisition timepoint. The files were then re-processed through the GUI-integrated segmentation network, which automatically generated binary masks of the vascular structures. The predicted masks were saved in a dedicated `"Pred"` subfolder with the naming structure `contrast_ST0XXYT_pred.nii`.

Quantitative microvascular features were automatically extracted using the GUI's *"Features Extraction"* module and exported as structured spreadsheet files.

A master JSON configuration file was used to maintain the correspondence between raw acquisitions, processed volumes, and automated segmentation outputs.

To quantitatively assess the effect of the optimization procedure, the microvascular features extracted from the optimized OCTA volumes were combined with the previously available radiomic descriptors and evaluated using the same longitudinal classification framework described in Section 2.4.1.

The resulting classification performance was compared with that obtained from the corresponding non-optimized dataset. This comparison was used to determine whether the correction of motion-corrupted acquisitions improved the discrimination among acquisition timepoints and increased the robustness of the longitudinal analysis.

2.3 OCT Morphological Processing and Feature Extraction Pipeline

A dedicated computational pipeline was developed in Python to reconstruct, preprocess, spatially realign, segment, and quantitatively analyze the acquired OCT/OCTA datasets. The software architecture was organized into two complementary Python modules:

- `oct_octa_pipeline.py`: the main processing module, containing the `OCT_OCTA_Processing` class. This class manages raw-data reconstruction, organization of repeated acquisitions, preliminary surface estimation, macroscopic volume realignment, artifact suppression, tissue-layer segmentation, visualization, and quantitative feature extraction.
- `Segmentation_algorithm_oct_octa.py`: a supporting image-processing module containing the routines used for three-dimensional preprocessing, skin-surface cost-function computation, graph-based surface segmentation, and Dermo-Epidermal Junction (DEJ) detection.

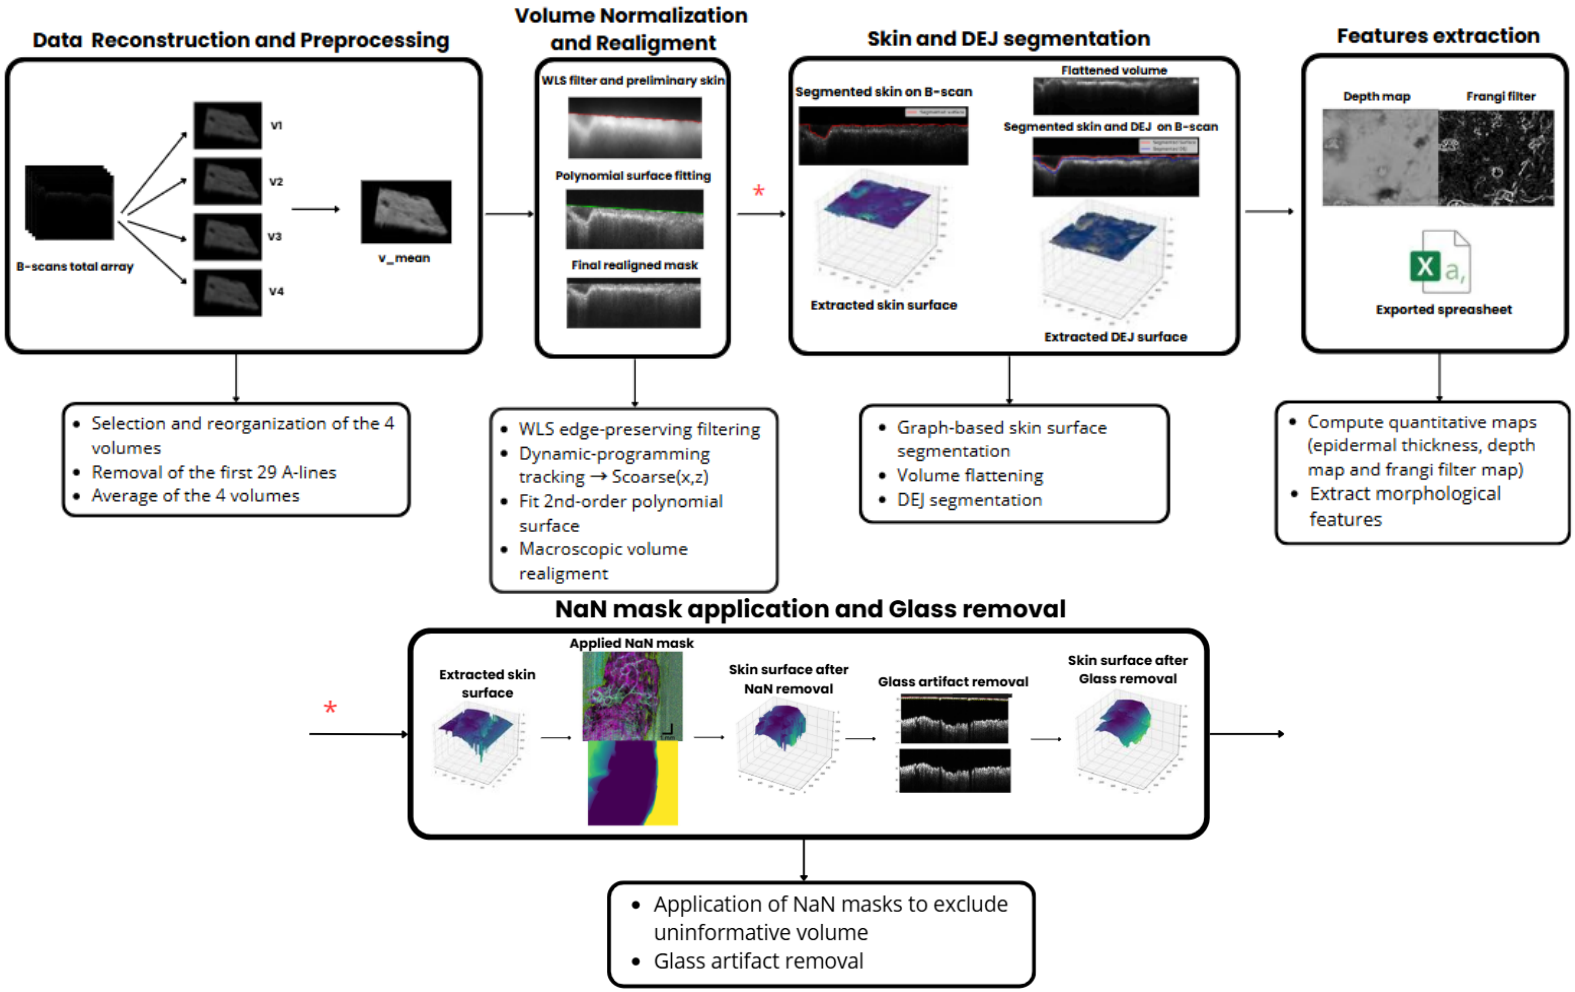


Figure 2.5: Overview of the OCT processing and feature extraction pipeline. The main workflow includes data reconstruction and preprocessing, volume normalization and realignment, skin-surface and DEJ segmentation, and feature extraction. The branch marked by (*) was applied only to acquisitions affected by partial or complete skin–probe detachment. For these volumes, a NaN exclusion mask was introduced to exclude unreliable regions from surface estimation, followed by glass-layer artifact removal, before proceeding with the main workflow.

2.3.1 Volume Preprocessing and Geometric Realignment

Separation of repeated OCT Acquisitions

At the beginning of the `intensity_based_OCT_OCTA()` procedure, the reconstructed dataset is flipped along the axial direction to restore the correct depth orientation. At each position, four consecutive B-scans were acquired before moving to the following position. The sequential dataset is therefore separated into four independent three-dimensional volumes, denoted as V_1, V_2, V_3, and V_4, corresponding to the four repeated acquisitions of the same spatial locations. Given the acquisition of 32 binary files containing 64 B-scans each, a total of 2048 B-scans were recorded for each acquisition. After separation of the four temporal repetitions, each reconstructed structural volume therefore contained 512 B-scan positions.

Removal of Lateral Acquisition Artifacts and Averaging of Repeated Acquisitions

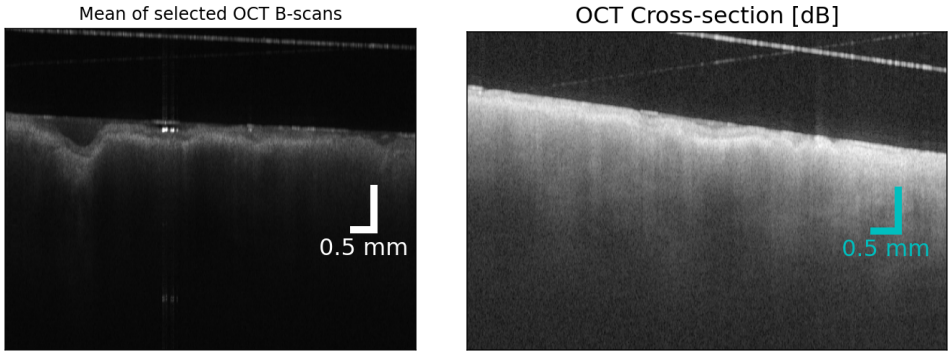
Before further processing, the first 29 A-lines of each B-scan are removed from all four reconstructed volumes. This region corresponds to the initial portion of the lateral scan and may contain galvo flyback or synchronization artifacts associated with the transition of the scanning system toward stable acquisition.

Following artifact removal, an averaged structural OCT volume, V_{mean}, is computed from the four repeated acquisitions.

Averaging the repeated acquisitions reduces temporal intensity fluctuations and acquisition noise while preserving the structural information consistently present across the four measurements. The resulting V_{mean} volume is subsequently used as the structural reference for preliminary surface tracking and macroscopic geometrical realignment.

Preliminary Surface Tracking by Edge-Preserving Filtering and Dynamic Programming

In contact-probe OCT imaging, the acquired skin surface may exhibit a substantial macroscopic inclination and curvature relative to the imaging plane (Figure 2.6). These large-scale geometrical variations must therefore be distinguished from local variations in skin topography, such as wrinkles, furrows, and lesion-related surface irregularities. If left uncorrected, the tilt could interfere with subsequent skin-layer segmentation and depth-dependent measurements. For this reason, a preliminary surface-tracking procedure was introduced before the refined anatomical segmentation. The purpose of this first step was not to obtain the final skin boundary used for quantitative analysis, but rather to generate a coarse geometrical reference, hereafter denoted as $S_{\text{coarse}}(x, y)$, from which the macroscopic orientation of the OCT volume could subsequently be estimated.



(a) **(b)**

Figure 2.6: Representative example of macroscopic skin-surface inclination in the original OCT volume. (a) Structural OCT B-scan showing the tilted skin surface. (b) Orthogonal cross-sectional view illustrating the corresponding surface curvature.

Weighted Least Squares Edge-Preserving Filtering Prior to preliminary surface tracking, the averaged structural OCT B-scans were processed using a Weighted Least Squares (WLS) edge-preserving filter [25]. To reduce computational requirements, the filtering operation was explicitly recalculated every fourth B-scan, and the corresponding preliminary surface estimate was propagated to the intermediate B-scans.

The OCT intensities were first clipped to positive values and transformed into the logarithmic domain:

$$L_{\log}(x, y) = \log\left[L(x, y)\right], \tag{2.1}$$

where $L(x, y)$ represents the structural OCT intensity.

Local forward and backward differences were calculated along both image directions. For a generic local logarithmic intensity difference d, the corresponding spatial weight was defined as:

$$w(d) = \frac{\lambda}{\epsilon + |d|^{\alpha}}, \tag{2.2}$$

where λ controls the overall regularization strength, α determines the sensitivity of the weighting function to local intensity gradients, and ϵ is a small positive constant introduced to prevent numerical instability. In the implemented algorithm,

$$\lambda = 50, \qquad \alpha = 2, \qquad \epsilon = 10^{-4}. \tag{2.3}$$

The magnitude of the weights is therefore larger in regions characterized by small local intensity differences and decreases in the presence of strong gradients.

Consequently, approximately homogeneous regions are subjected to stronger spatial regularization, whereas prominent intensity transitions are comparatively preserved.

The resulting weights were used to construct a sparse linear system (Ax=b), in which the diagonal coefficients account for the cumulative contribution of the four neighboring pixels, while the corresponding off-diagonal coefficients are given by the negative directional weights. Boundary connections were explicitly removed to avoid interactions with non-existing or non-adjacent pixels. The system was then solved using a sparse linear solver and the resulting vector reshaped to the original B-scan dimensions, providing an edge-preserving smoothed image for the subsequent processing steps.

Dynamic-programming surface tracking The WLS-filtered B-scan was subsequently processed using a dynamic-programming algorithm to identify a continuous low-cost path across the image.

The absolute horizontal and vertical image gradients were first calculated and independently normalized between 0 and 1, resulting in $G_{x,N}$ and $G_{y,N}$. A local cost was then assigned to each pixel:

$$W(y,x) = [1 - G_{x,N}(y,x)] + [1 - G_{y,N}(y,x)] + \sigma, \tag{2.4}$$

where $\sigma = 0.01$ is a small positive constant ensuring a non-zero cost.

Pixels characterized by strong intensity gradients therefore exhibit lower costs and are preferentially selected as candidate positions for the superficial tissue interface.

Starting from the first image column, the cumulative minimum cost was propagated from left to right according to:

$$C(y,x) = W(y,x) + \min_{\Delta y \in [-m,m]} C(y + \Delta y, x - 1), \tag{2.5}$$

where m represents the maximum permitted axial displacement between two adjacent A-lines. In the implemented pipeline,

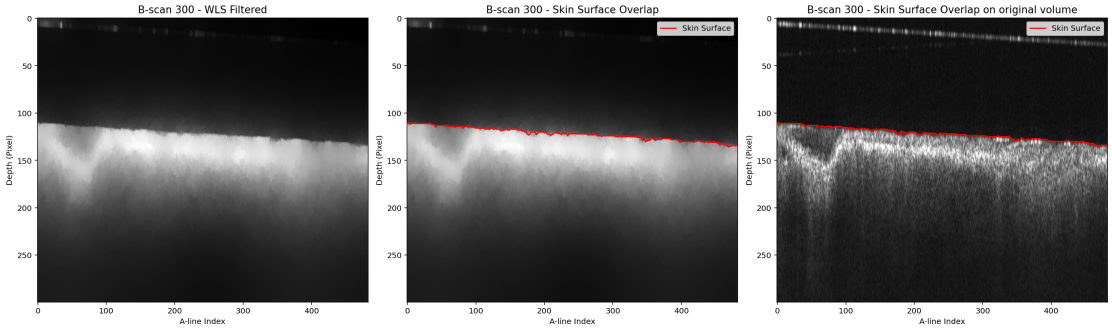
$$m = 2 \text{ pixels}. \tag{2.6}$$

The ± 2-pixel constraint prevents abrupt variations of the detected path while allowing gradual changes in surface position. After reaching the final image column, backtracking through the stored minimum-cost predecessors provides the complete preliminary contour.

The individual contours were assembled across the volume to obtain the two-dimensional coarse surface map:

$$S_{\text{coarse}}(x,z). \tag{2.7}$$

S_{coarse} represents an approximate geometrical representation of the skin.



(a) (b) (c)

Figure 2.7: Preliminary skin-surface tracking in a representative structural OCT B-scan. (a) B-scan after WLS edge-preserving filtering. (b) Preliminary surface detected by dynamic programming and superimposed on the filtered B-scan. (c) Same preliminary surface superimposed on the corresponding original B-scan.

OCTA-Derived Exclusion of Unreliable Preliminary Surface Regions

Contact-probe acquisitions may present different optical conditions depending on the degree of contact between the skin and the probe window:

- Under firm contact conditions, the tissue remains optically coupled to the probe window and the superficial structural OCT signal can be reliably associated with the skin surface. Optical coupling is maintained, and the primary reflective interface accurately corresponds to the stratum corneum.
- In the opposite condition, local detachment of the skin from the probe may introduce an air gap between the glass window and the tissue. The tissue detaches from the window. The refractive index mismatch ($n_{\text{glass}} \approx 1.52 \rightarrow n_{\text{air}} \approx 1.0$) causes a strong specular reflection at the glass surface, while structural light reaching the detached skin below undergoes severe attenuation. Consequently, the superficial glass reflection may become more prominent than the actual skin boundary in the structural OCT image.

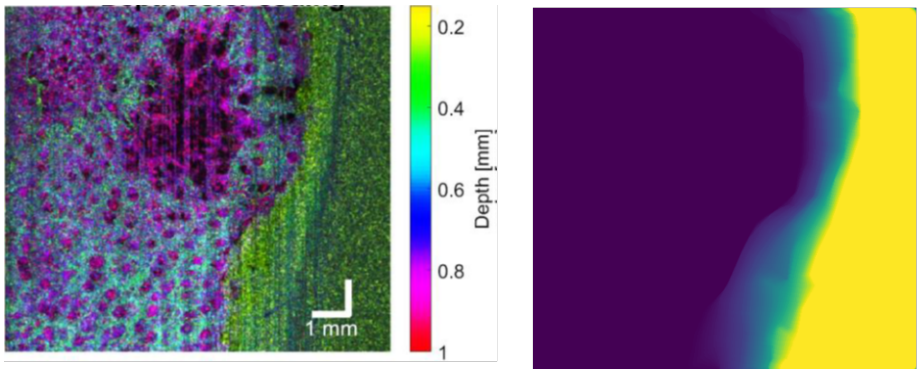
A critical limitation of the preliminary surface-tracking procedure arises in regions where the skin locally detaches from the probe window. In these areas, the dynamic-programming algorithm may preferentially track the strong specular reflection generated at the glass-air interface rather than the actual skin boundary. This misidentification introduces severe inaccuracies in the coarse surface map S_{coarse}, compromising its ability to faithfully represent the underlying tissue morphology.

This artifact is particularly pronounced in anatomically challenging regions, such as the nose, ears, or areas characterized by high surface curvature and small contact areas, where consistent mechanical coupling with the probe window cannot be reliably maintained.

Regions affected by probe detachment exhibit a pronounced loss of functional vascular signal in the corresponding OCTA data, due to the severe optical attenuation of the beam propagating across the air gap.

In previous work, these compromised zones were manually segmented on the OCTA volumes to generate 3D binary exclusion masks. In the present pipeline, these prior annotations were exploited to construct selective exclusion masks.

Only masks specifically associated with tissue detachment were retained, while masks generated for other sources of signal degradation (such as severe motion artifacts) were excluded. An example of exclusion mask alongside their corresponding OCT/OCTA volumes are illustrated in Figure 2.8. The volumes requiring this targeted correction present the path of the mask in the json file (`patient_files_maskoct.json`). The volumes which required the mask application are detailed in Table 2.5.



(a) **(b)**

Figure 2.8: Representative correspondence between an uninformative OCTA region and the associated exclusion mask. (a) Depth-color-coded en-face OCTA projection showing the region affected by signal loss. (b) Corresponding exclusion mask, where violet indicates valid tissue regions and yellow identifies areas affected by tissue detachment.

Table 2.5: Overview of the selected OCT/OCTA volumes and corresponding acquisition timepoints included in the analysis.

Volume ID	Selected Timepoints
ST004.1	post, 3M, 6M, healthy
ST004.2	pre, 3M, 6M
ST008.1	3M
ST010.1	pre, post, 3M
ST014.1	pre, post, 3M, 6M, healthy
ST014.2	pre, post, 3M, 6M
ST015.1	pre, healthy

The locations identified by the exclusion mask were then assigned `NaN` values in the coarse surface map to prevent unreliable surface coordinates from contributing to the subsequent estimation of the macroscopic tissue orientation:

$$S_{\text{coarse}}(x, y) = \text{NaN} \qquad \text{if} \qquad M_{\text{excl}}(x, y) = 1. \tag{2.8}$$

The resulting masked surface can therefore be expressed conceptually as:

$$S^{*}_{\text{coarse}}(x, y) = \begin{cases} S_{\text{coarse}}(x, y), & \text{reliable region,} \\ \text{NaN}, & \text{excluded region.} \end{cases} \tag{2.9}$$

Only the valid coordinates of S^{*}_{coarse} were used for polynomial estimation of the macroscopic surface trend. In this way, regions in which the preliminary algorithm could erroneously follow the probe-window reflection were prevented from biasing the geometrical realignment of the OCT volume. However, for the acquisitions listed in Table 2.5, the partial or complete loss of contact between the probe window and the skin produced particularly pronounced glass-layer reflections. In these cases, the exclusion mask alone was not sufficient to prevent the residual glass signal from interfering with the subsequent skin-surface segmentation, particularly near the boundaries of the valid imaging region. Therefore, an additional glass-layer artifact removal procedure was developed and applied to these selected volumes before refined surface segmentation.

Macroscopic Surface Modeling and Volume Flattening

Directly flattening the OCT volume according to the preliminary surface map would remove not only the undesired macroscopic inclination, but also physiologically meaningful local topographical variations. The preliminary surface is therefore used exclusively to estimate the large-scale geometrical trend of the tissue.

To correct the global tilt without attenuating local surface irregularities (wrinkles, epidermal furrows, open lesions), a smooth second-order polynomial surface is fitted to the valid coordinates of the preliminary surface map:

$$S_{poly}(x, z) = c_1x^2 + c_2z^2 + c_3xz + c_4x + c_5z + c_6 \tag{2.10}$$

The polynomial coefficients are estimated by least-squares fitting of the valid coordinates of the preliminary surface map.

The fitted polynomial surface was subsequently used to correct the macroscopic orientation of the OCT volume. For each lateral position (x, z), the corresponding A-line was axially shifted so that the polynomial surface $S_{\text{poly}}(x, z)$ was aligned to a common reference depth of 20 pixels.

The same spatial transformation was applied to all four repeated OCT volumes. Voxels shifted outside the original axial range were assigned zero intensity.

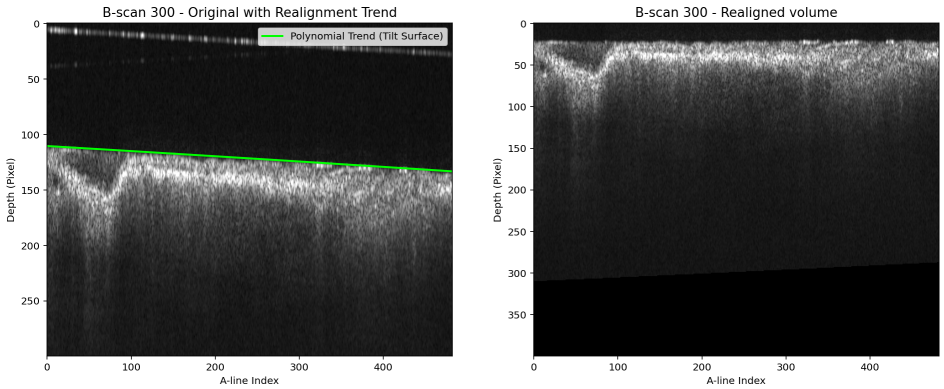


Figure 2.9: Representative structural OCT B-scan before and after macroscopic surface realignment, showing the correction of large-scale tissue inclination while preserving local skin-surface topography.

Glass Layer Artifact Removal

The glass layer, where the previous steps were not capable of removing it, appears as a more or less defined straight line above the surface or in partial contact with it, with a thickness of a few pixels (2.10). Glass-layer artifact removal was applied only to acquisitions for which an exclusion mask was provided. The volume intensities were first rescaled between the 2nd and 98th percentiles to reduce the influence of extreme intensity values. A global intensity threshold was then defined as

$$T_{\text{glass}} = 0.1\, I_{\max}, \tag{2.11}$$

where $I_{\max}$ represents the maximum intensity of the rescaled volume.

For each lateral position (x, z), the first axial location (from the top) exceeding this threshold was identified as the approximate glass-interface position:

$$y_{\text{glass}}(x, z) = \min \{y : I(y, x, z) > T_{\text{glass}}\} . \tag{2.12}$$

A 5×5 median filter was subsequently applied to the resulting two-dimensional interface map to reduce isolated detection errors while preserving its overall spatial distribution. A band extending from 50 pixels above to 9 pixels below the detected interface was then removed:

$$y \in [y_{\text{glass}}(x, z) - 50,\ y_{\text{glass}}(x, z) + 9] . \tag{2.13}$$

Voxel intensities within this interval were set to zero. This asymmetric removal window was selected to suppress the highly reflective glass-related structure and its surrounding signal while limiting the removal of deeper tissue information.

Regions previously identified as unreliable through the exclusion mask were not used for subsequent quantitative segmentation and analysis. The resulting artifact-suppressed volume was then provided as input to the refined skin-surface segmentation stage.

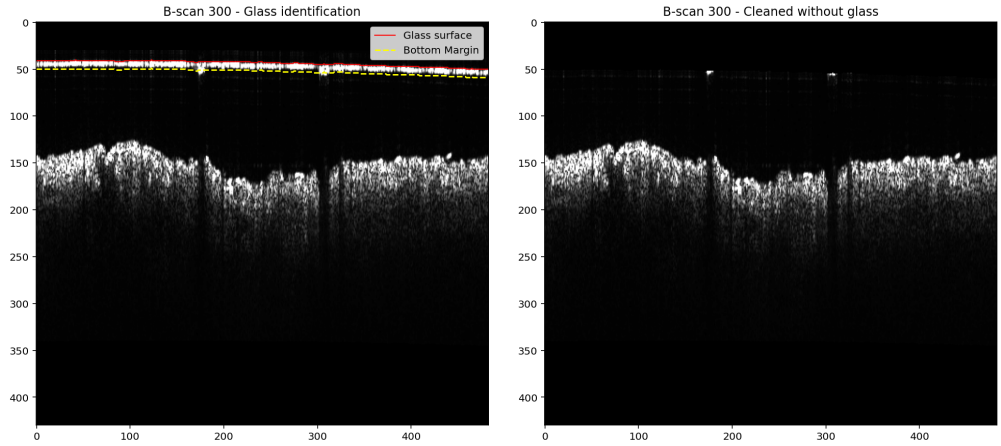


Figure 2.10: Glass artifact suppression. Left: Identification of the high-intensity glass reflection interface and defined masking margins. Right: Cleaned volume with the glass reflection suppressed.

2.3.2 Skin Surface and Dermal-Epidermal Junction Segmentation

Graph-Based Skin Surface Segmentation

Following macroscopic realignment and, if required, glass-layer artifact removal, the structural OCT volume was further processed to obtain a refined segmentation of the skin surface. An adapted two-dimensional implementation of the graph-based framework proposed by [14] was developed, as schematically illustrated in Figure 2.11.

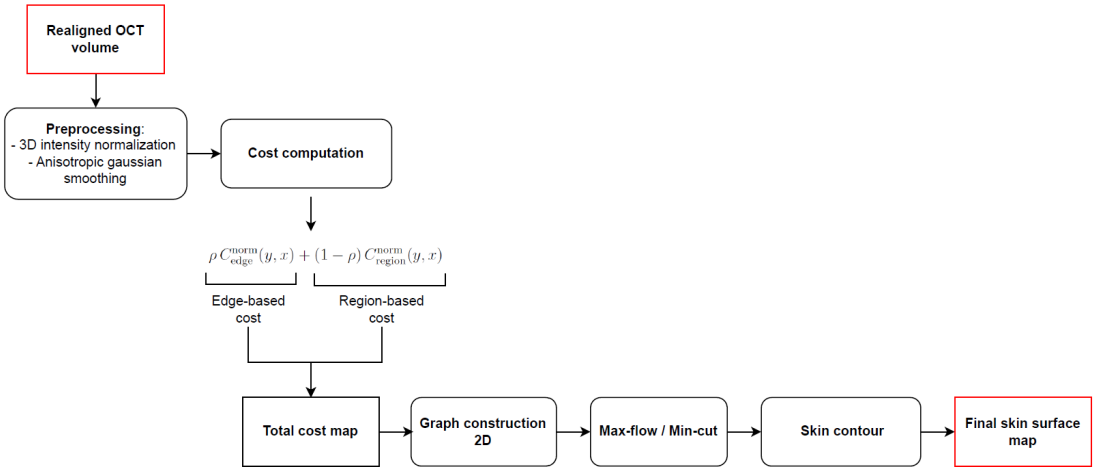


Figure 2.11: Schematic representation of the developed graph-based skin-surface segmentation pipeline. For each B-scan, edge- and region-based information is combined into a cost map, which is converted into a directed graph and optimized through a max-flow/min-cut approach to identify the skin surface.

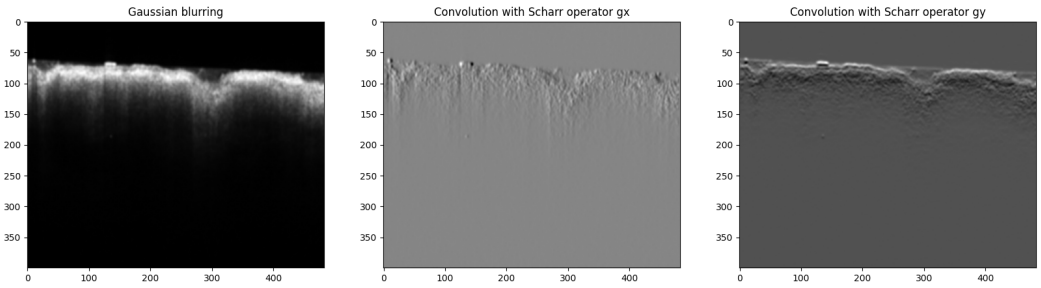
The realigned mean OCT volume was first converted to an 8-bit representation and contrast-rescaled using the 2nd and 98th intensity percentiles. An anisotropic Gaussian filter was subsequently applied to suppress high-frequency noise while preserving the main structural transitions.

The preprocessed volume was then analyzed B-scan by B-scan. For each B-scan, a cost function combining edge- and region-based information was constructed.

The edge-based contribution was obtained using Scharr operators to estimate the horizontal and vertical intensity gradients, g_x and g_y. The corresponding gradient orientation was computed as

$$\phi_z(x,y) = \arctan\left(\frac{g_y}{g_x}\right). \tag{2.14}$$

and converted to degrees.



(a) (b) (c)

Figure 2.12: Gradient computation for the edge-based cost. (a) Representative preprocessed structural OCT B-scan. (b) Horizontal gradient component g_x obtained using the Scharr operator. (c) Vertical gradient component g_y obtained using the Scharr operator.

An orientation-dependent weighting function $p(\phi_z)$ was introduced to favor dark-to-bright transition at the skin surface:

$$p(\phi_z) = \begin{cases} 1, & 0° \leq \phi_z \leq 180°, \\ 1 - \dfrac{\phi_z - 180°}{90°}, & 180° < \phi_z < 270°, \\ 0, & \phi_z = 270°, \\ \dfrac{\phi_z - 270°}{90°}, & 270° < \phi_z < 360°. \end{cases} \tag{2.15}$$

Thus, orientations associated with the expected superficial transition were assigned the maximum weight, whereas the weight linearly decreased toward the unfavorable orientation at 270°.

In the developed implementation, the orientation information was additionally combined with the local gradient magnitude, which generated a cost map as follows (Figure 2.13):

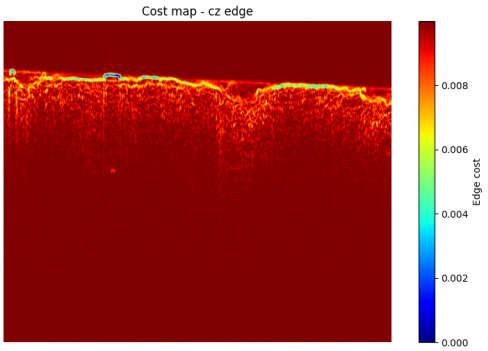


Figure 2.13: Representative edge-based cost map obtained by combining the normalized Scharr gradient magnitude with the orientation-dependent weighting function. Low-cost regions preferentially correspond to strong dark-to-bright transitions compatible with the expected skin-surface boundary.

The gradient magnitude was calculated from the Scharr components as

$$M(x,y) = \sqrt{g_x(x,y)^2 + g_y(x,y)^2} \tag{2.16}$$

and independently normalized to the range [0,1]:

$$M_{\text{norm}}(x,y) = \frac{M(x,y) - M_{\text{min}}}{M_{\text{max}} - M_{\text{min}}}. \tag{2.17}$$

The edge-based cost was then defined as

$$C_{\text{edge}}(x,y) = 1 - p(\phi_z(x,y))\, M_{\text{norm}}(x,y). \tag{2.18}$$

This formulation simultaneously accounts for the orientation and strength of the local intensity transition. Consequently, a strong gradient with an orientation compatible with the expected skin boundary produces a value of $p(\phi_z)M_{\text{norm}}$ close to 1 and therefore a low edge cost, which will be preferentially selected. Conversely, weak gradients or gradients with unfavorable orientations result in higher cost values.

In parallel, a region-based contribution was computed from the structural OCT intensity. Representatives B-scan intensity maps were visualized to identify the threshold which allowed distinguishing the low-intensity region above the sample from the brighter, optically scattering tissue. To this purpose, a fixed intensity threshold $\theta = 100$ was applied, generating the binary brightness map

$$T(y,x) = \begin{cases} 1, & I(y,x) \geq 100, \\ 0, & I(y,x) < 100. \end{cases} \tag{2.19}$$

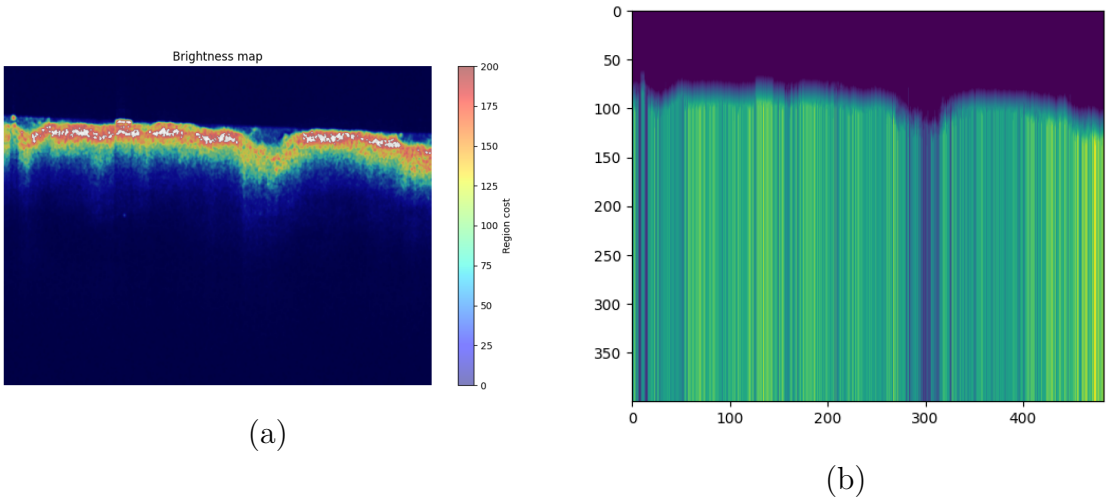
For each A-line, the detected bright pixels were then cumulatively integrated along the axial direction:

$$B(y,x) = \sum_{k=0}^{y} T(k,x). \tag{2.20}$$

Consequently, pixels located above the tissue surface, where few or no intensities exceed the selected threshold, are associated with low cumulative values. Once the superficial tissue boundary is crossed, the increasing number of bright pixels produces a progressive increase in $B(y,x)$ along the axial direction. This generates a region-based map encoding the transition from the dark background above the sample toward the brighter scattering tissue.

The cumulative map was subsequently normalized to the range [0,1] to obtain the region-based cost:

$$C_{\text{region}}(y,x) = \frac{B(y,x) - B_{\min}}{B_{\max} - B_{\min}}. \tag{2.21}$$



(a) (b)

Figure 2.14: Representative construction of the region-based cost. (a) Structural OCT intensity map. Pixels exceeding the selected intensity threshold ($\theta = 100$) were identified as bright tissue pixels. (b) Cumulative bright-pixel map obtained by integrating the thresholded binary values along the axial direction for each A-line. The resulting map provides the region-based contribution used for skin-surface segmentation.

The independently normalized edge and region terms were finally combined into a single cost map:

$$C_{\text{total}}(y,x) = \rho\, C_{\text{edge}}^{\text{norm}}(y,x) + (1-\rho)\, C_{\text{region}}^{\text{norm}}(y,x), \tag{2.22}$$

where $\rho = 0.5$ was used to assign equal contribution to the edge- and region-based information.

For each B-scan, the resulting cost map was converted into a directed graph $G = (V, E)$, in which each image pixel corresponded to a graph node. Terminal edges connected each node to the source and sink according to the vertical differences between adjacent values of C_{total}.

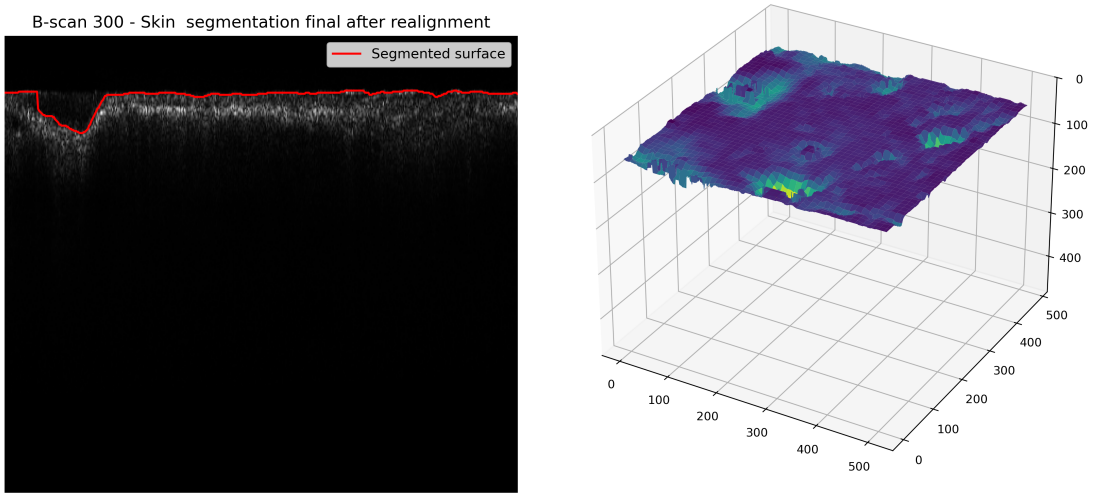
The minimum source–sink cut was then obtained through max-flow/min-cut optimization. A contour $y_{\text{skin}}(x)$ for was obtained for each B-scan.

The contours extracted from all B-scans were assembled into the two-dimensional surface map

$$Surface_{\text{skin}}(x,z), \tag{2.23}$$

A two-dimensional median filter was subsequently applied to the assembled surface map to reduce isolated segmentation irregularities and improve spatial consistency.

The resulting skin-surface segmentation is illustrated in Figure 2.15.



(a) (b)

Figure 2.15: Refined skin-surface segmentation. (a) Representative structural OCT B-scan with the segmented skin surface superimposed in red. (b) Three-dimensional representation of the resulting skin-surface map.

Dermo-Epidermal Junction Detection

Following refined skin-surface segmentation, the Dermo-Epidermal Junction (DEJ) was identified from the structural OCT volume. Since the DEJ generally produces a less pronounced intensity transition than the superficial skin boundary, its detection was performed after spatially normalizing the tissue with respect to the previously segmented skin surface.

For each B-scan, every A-line was axially shifted according to the corresponding coordinate of $Surface_{\text{skin}}(x, z)$. This transformation generated a flattened volume in which the underlying tissue layers could be analyzed independently of residual local skin topography or residual background noise.

Prior to DEJ detection, the flattened volume was normalized between its minimum and maximum intensity values and subsequently smoothed using an anisotropic three-dimensional Gaussian filter ($\sigma_{\text{depth}} = 1.5, \sigma_{\text{A-line}} = 1, \sigma_{\text{B-scan}} = 5$).

Contrast enhancement was then performed using Contrast-Limited Adaptive Histogram Equalization (CLAHE) with a three-dimensional kernel size of 8x5x5 voxels.

DEJ detection was subsequently performed independently on each B-scan. For every A-line position x, a local search window was defined within the first 30 axial pixels below the flattened skin surface and across a small neighborhood of adjacent A-lines. The structural intensities within this lateral window were summed to generate a one-dimensional axial intensity profile:

$$P_x(y) = \sum_{k=x-w}^{x+w} I_{\text{CLAHE}}(y, k), \tag{2.24}$$

where I_{CLAHE} represents the preprocessed structural OCT B-scan and the lateral neighborhood was centered on the A-line under analysis.

Local maxima in $P_x(y)$ were detected. When at least two significant peaks were identified, the first two peaks, denoted as p_1 and p_2, were considered. These peaks represent the main superficial intensity structures detected within the local tissue profile.

The DEJ candidate was defined as the axial position corresponding to the minimum intensity value within the interval separating the two peaks (2.16):

$$y_{\text{DEJ}}(x) = \arg\min_{y \in [p_1, p_2]} P_x(y). \tag{2.25}$$

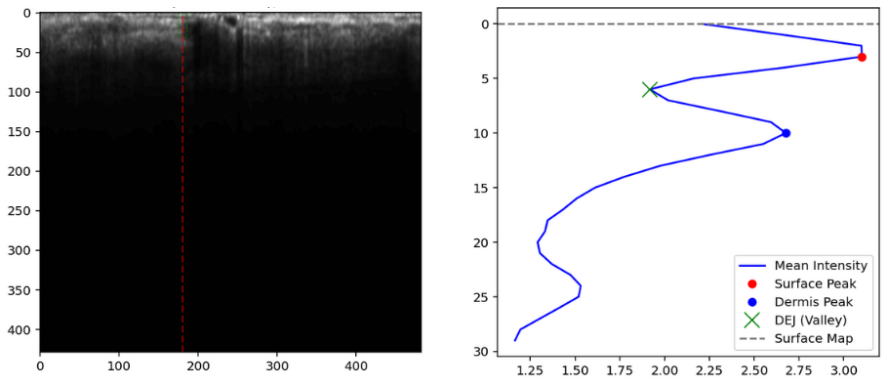


Figure 2.16: Representative DEJ detection from the local axial intensity profile. The first two significant intensity peaks (p_1 and p_2) are identified within the search window, and the DEJ candidate is assigned to the minimum-intensity position between them.

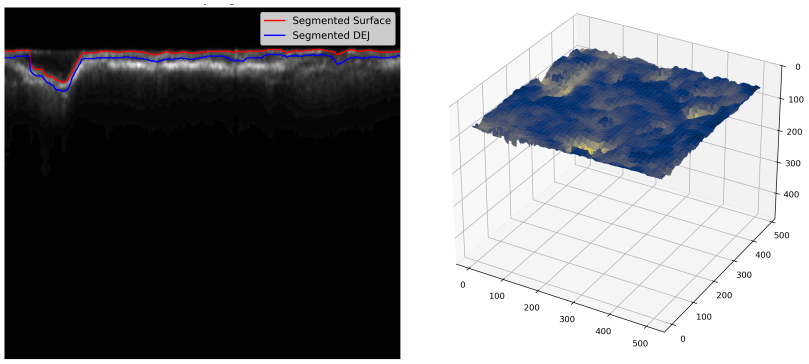
This approach exploits the characteristic reduction in OCT intensity occurring between the highly scattering superficial epidermal structures and the underlying dermal tissue.

To limit abrupt variations between neighboring A-lines, the detected DEJ coordinate was constrained with respect to the preceding estimate, restricting the maximum allowed axial variation to 2 pixels. The resulting DEJ contour was subsequently smoothed using a one-dimensional median filter.

The DEJ contours obtained from all B-scans were then assembled into a two-dimensional DEJ depth map.

Because DEJ detection was performed in the skin-flattened coordinate system, the resulting map represents the local axial distance between the superficial skin boundary and the detected DEJ. For visualization in the original spatial reference system, the skin-surface coordinates were added back to the DEJ map.

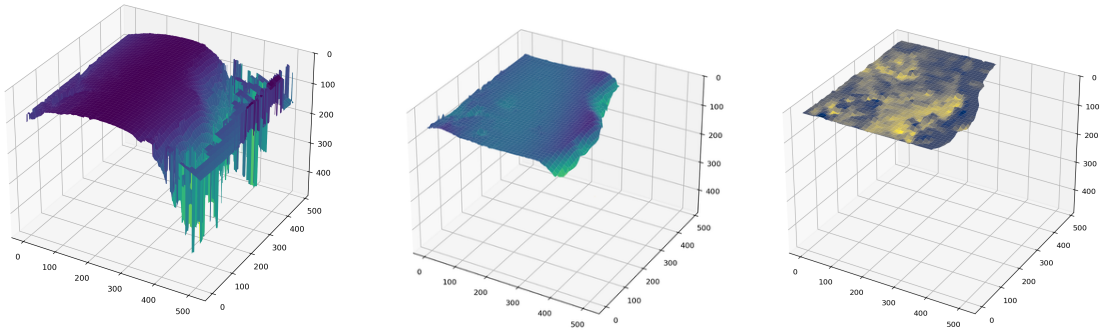
The final segmentation of the superficial skin boundary and DEJ is shown in Figure 2.17.



(a) **(b)**

Figure 2.17: Skin-surface and Dermo-Epidermal Junction segmentation. (a) Representative structural OCT B-scan showing the segmented skin surface and DEJ. (b) Three-dimensional representation of the detected DEJ surface.

For patients requiring a NaN exclusion mask, the identified invalid regions were removed from the segmented skin and DEJ surface maps. Consequently, surface coordinates within the masked regions were excluded from subsequent analysis, resulting in the maps representations illustrated in Figure 2.18.



(a) **(b)** **(c)**

Figure 2.18: Representative effect of the NaN exclusion mask on the segmented skin and Dermo-Epidermal Junction (DEJ) surface maps. (a) Three-dimensional skin-surface map before application of the exclusion mask. (b) Corresponding skin-surface map after removal of the masked regions. (c) Corresponding DEJ surface map after removal of the masked regions.

2.3.3 Quantitative Feature Extraction

From the segmented boundaries (y_{skin} and y_{DEJ}), a set of quantitative morphometric features is extracted to characterize tissue microarchitecture.

Epidermal Thickness

Epidermal thickness $T_{epi}(x, z)$ is defined as the vertical distance between the skin surface and the DEJ. Pixel distances are converted to micrometers (μm) using the axial pixel dimension $dz = 0.0137$ mm and the tissue refractive index $n = 1.43$:

$$T_{epi}(x, z) = \frac{(y_{DEJ}(x, z) - y_{skin}(x, z)) \times dz \times 1000}{n} \quad [\mu\text{m}] \tag{2.26}$$

The mean epidermal thickness ($\bar{T}_{epi}$) and its standard deviation (σ_T) are computed across all valid non-NaN positions:

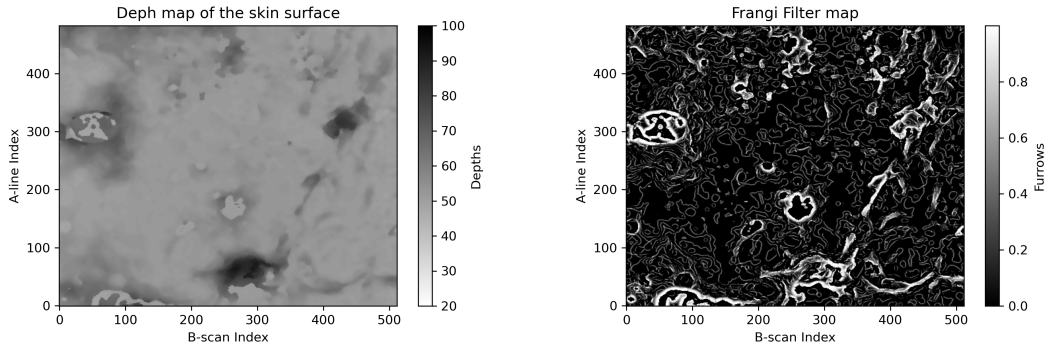
$$\bar{T}_{epi} = \frac{1}{N_{valid}} \sum_{(x,z)\in\Omega} T_{epi}(x, z) \tag{2.27}$$

$$\sigma_T = \sqrt{\frac{1}{N_{valid}} \sum_{(x,z)\in\Omega} \left(T_{epi}(x, z) - \bar{T}_{epi}\right)^2} \tag{2.28}$$

Surface Topography and Roughness Analysis

Skin surface roughness was quantified from the previously extracted skin-surface map. Following the approach proposed by [14], the detected surface coordinates were first represented as a two-dimensional depth map, $Surface_{\text{skin}}(x, z)$, in which each pixel value corresponds to the axial position of the skin surface at the respective A-line and B-scan coordinates. Therefore, local variations in the map represent variations in skin-surface topography.

A two-dimensional Frangi filter was directly applied to the skin depth map. The filter was used to detect structures across multiple spatial scales (σ = 1,2,3 pixels), using $\alpha = 0.5$, $\beta = 0.5$, and $\gamma = 0.2$. The resulting Frangi response assigns higher values to ridge-like structures in the depth map, thereby providing a spatial representation of the detected surface furrows. A threshold of 0.5 was subsequently applied to the filter response to define the pixels classified as furrows. Representative skin depth and Frangi-response maps are shown in Figure 2.19.



(a) **(b)**

Figure 2.19: Skin-surface roughness analysis. (a) Depth map obtained from the segmented skin surface, where pixel values represent the detected axial surface position. (b) Corresponding Frangi-filter response used to identify surface furrows.

Three topographical metrics are extracted:

1. **Furrow Density ($F_{density}$):** The ratio of furrow pixels (N_{furrow}) to the total skin area (N_{total}):
$$F_{density} = \frac{N_{furrow}}{N_{total}} \tag{2.29}$$

2. **Furrow Depth (F_{depth}):** The mean axial deviation of furrow locations from the baseline surface elevation ($\bar{Z}_{skin}$):
$$F_{depth} = \frac{1}{N_{furrow}} \sum_{(x,z)\in\text{furrows}} \left| Z_{skin}(x,z) - \bar{Z}_{skin} \right| \quad \text{[pixels]} \tag{2.30}$$
F_{depth} is subsequently converted to physical units (μm).

3. **Maximum Furrow Depth (D_{max}):** The 99th percentile depth deviation within furrow regions, quantifying the maximum vertical extent of surface depressions:
$$D_{99} = \text{Percentile}_{99}\left(\left\{ Z_{\text{skin}}(x,z) - \bar{Z}_{\text{skin}} \,\middle|\, (x,z) \in \Omega_{\text{furrow}} \right\}\right) \quad \text{[pixels]}, \tag{2.31}$$
where Ω_{furrow} denotes the set of pixels identified as furrows from the thresholded Frangi response. The resulting value was converted to physical units according to
$$D_{99,\mu\text{m}} = \frac{D_{99}\, d_z\, 1000}{n}, \tag{2.32}$$

where $d_z = 0.0137$ mm is the axial pixel size and $n = 1.43$ is the refractive index of the skin.

All extracted metrics are exported to a structured spreadsheet for subsequent statistical analysis.

2.4 Statistical Analysis

The extracted quantitative imaging biomarkers were subsequently analyzed to investigate their association with treatment-related temporal changes and clinically assess radiation-induced toxicity.

2.4.1 Longitudinal Classification

A longitudinal classification framework was employed to assess whether the extracted quantitative imaging features could discriminate among different acquisition timepoints and, consequently, capture treatment-related temporal changes. The same classification framework was used for two complementary analyses.

First, as part of the data optimization procedure described in Section 2.2, longitudinal classification was used to quantitatively evaluate the effect of OCTA volume optimization. Classification performance obtained from the optimized dataset was compared with that obtained from the corresponding non-optimized dataset, providing a quantitative measure of whether artifact correction improved the temporal discriminative information contained in the extracted OCTA and radiomic features.

Second, the longitudinal classification analysis was used to investigate the contribution of the OCT-derived morphological biomarkers introduced in the present work. For this purpose, two feature representations were considered: a reference feature set containing OCTA-derived vascular parameters and radiomic descriptors, and an extended feature set additionally including the OCT-derived morphological features. Comparison of the corresponding classification performances was used to assess whether the addition of structural OCT information improved discrimination among longitudinal acquisition stages.

For both analyses, the acquisition timepoint was used as the multiclass target variable, comprising baseline (PRE), immediately post-treatment (POST), 3-month follow-up (3M), 6-month follow-up (6M), and healthy skin (HEALTHY), where available.

The classification framework was implemented in Python following the methodology described by [16].

A scheme describing the followed steps is represented in Figure 2.20.

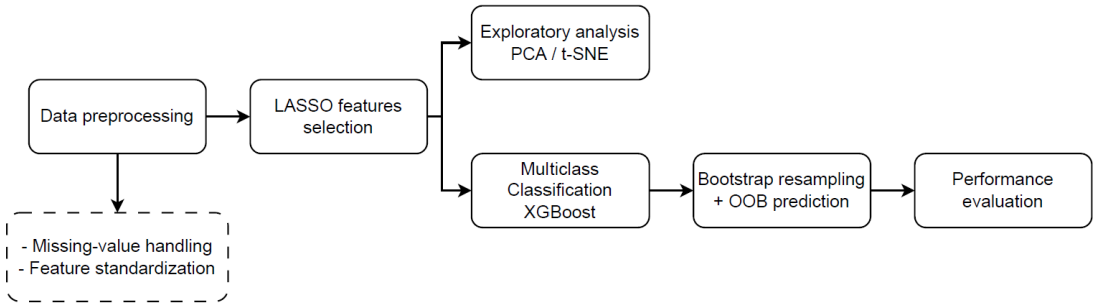


Figure 2.20: Schematic representation of the longitudinal classification analysis workflow. The process includes data preprocessing, feature selection, dimensionality reduction, and multiclass classification using an XGBoost model.

Before classification, features with excessive missing information were excluded, while remaining missing values were replaced by the median value of the corresponding feature. Quantitative variables were standardized to ensure comparable feature scales. Feature dimensionality was subsequently reduced using LASSO regression, retaining variables associated with non-zero regression coefficients. The selected quantitative biomarkers were further explored through Principal Component Analysis (PCA) and t-distributed Stochastic Neighbor Embedding (t-SNE) to investigate whether their distribution already had different patterns across timepoints. Multiclass classification was subsequently performed using an Extreme Gradient Boosting (XGBoost) classifier. Model performance was evaluated through bootstrap resampling with out-of-bag (OOB) prediction. At each iteration, observations sampled with replacement were used for model training, whereas observations not selected in the corresponding bootstrap sample constituted the OOB evaluation set. For observations evaluated in multiple bootstrap iterations, predicted class probabilities were averaged across all available OOB predictions, providing an aggregated prediction for each acquisition.

2.4.2 Clinical Toxicity Scoring

Radiation-induced skin toxicity was clinically assessed following the CTCAE v5.0 criteria, which provide a standardized framework for grading adverse events in oncology trials. Each patient was evaluated at multiple timepoints post-radiotherapy, and the severity of skin reactions was recorded on a scale from Grade 1 (mild toxicity) to Grade 5 (severe toxicity). The clinical assessments were performed by trained clinicians and documented in the patients' medical records. The registered toxicity scores were ranging from 0 to 3. To facilitate the analysis, these scores were subsequently grouped into two classes: Class 0 (representing no or mild toxicity, corresponding to Grades 0 and 1) and Class 1 (representing moderate to severe

toxicity, corresponding to Grades 2 and 3). This binary classification allowed for a more robust statistical analysis of the association between imaging biomarkers and clinically relevant toxicity outcomes. In Table 2.6 the distribution of toxicity scores across the patient cohort is summarized, highlighting the number of patients and lesions corresponding to each toxicity class.

Table 2.6: Patient and lesion distribution across toxicity classes. The table presents the number of patients and total lesions corresponding to each toxicity grade, as well as the aggregated counts for the two defined classes.

Grade of Toxicity	**Patients** n **(%)**	**Total Lesions** n
Class 0 (Toxicity 0 – 1)	**12 (63.2%)**	**30**
↪ Toxicity 0	0 (0%)	0
↪ Toxicity 1	12 (63.2%)	30
Class 1 (Toxicity 2 – 3)	**7 (36.8%)**	**15**
↪ Toxicity 2	6 (31.6%)	11
↪ Toxicity 3	1 (5.3%)	4
Total	**19 (100%)**	**45**

2.4.3 Toxicity Discrimination Model

To investigate the association between quantitative imaging biomarkers and radiation-induced toxicity, three feature representations were considered.

- **PRE:** Baseline features extracted from the pre-treatment acquisition.
- **PRE + ΔPOST:** Baseline features combined with their longitudinal variations registered at the post-treatment acquisition.
- **PRE + Δ3M:** Baseline features combined with their longitudinal variations registered at the 3-month follow-up acquisition.

Feature differences were calculated at lesion level as the follow-up value minus the corresponding baseline value. In mixed datasets (PRE + ΔPOST and PRE + Δ3M), only lesions with both baseline and follow-up acquisitions were included in the analysis (some patients had missing data). The resulting composition of the three different feature sets is summarized in the Table below (Table 2.7):

Table 2.7: Overall dataset composition detailing total patients, Toxicity 0 and Toxicity 1 counts, and total lesions across PRE, PRE + ΔPOST, and PRE + Δ3M cohorts.

Dataset	**Total Patients** n	**Class 0** n (%)	**Class 1** n (%)	**Total Lesions** n
PRE	19	12 (63.2%)	7 (36.8%)	43
PRE + ΔPOST	17	11 (64.7%)	6 (35.3%)	39
PRE + Δ3M	17	12 (70.6%)	5 (29.4%)	37

This design allowed both baseline tissue characteristics and early treatment-related changes to be evaluated as potential predictors of toxicity.

Because multiple lesions could belong to the same patient, model validation was performed using patient-level grouping. A regularized logistic regression model was used for this statistical analysis, as illustrated in the pipeline scheme (Figure 2.21). Data were split into training and test folds using a patient-level GroupKFold cross-validation scheme. This ensured that all lesions from a given patient were assigned exclusively to either the training or the test set in each iteration.

Data preprocessing, comprising missing-value imputation, feature standardization, and feature selection, was performed exclusively within each outer training fold. Feature selection was conducted using a stability selection procedure: repeated patient-level bootstrap resampling of the training fold was performed, fitting an L1-regularized logistic regression model to each resampled subset. The frequency with which each feature received a non-zero coefficient served as a measure of stability, and the most consistently selected features were retained for model training.

Three regularized logistic modeling strategies were evaluated: a Logistic Stable classifier, Ridge logistic regression, and ElasticNet logistic regression. While the Logistic Stable model relied solely on the selected stable features, Ridge and ElasticNet required hyperparameter tuning. This tuning was performed via an inner patient-grouped cross-validation procedure, where each outer training fold was further split into inner training and validation folds using patient-level grouping. The hyperparameter configuration yielding the best performance on the inner validation folds was selected for final model fitting.

Finally, the optimized models generated predictions exclusively for observations in the held-out outer test folds. Out-of-fold predictions across all five folds were pooled to evaluate overall model performance for each of the three approaches.

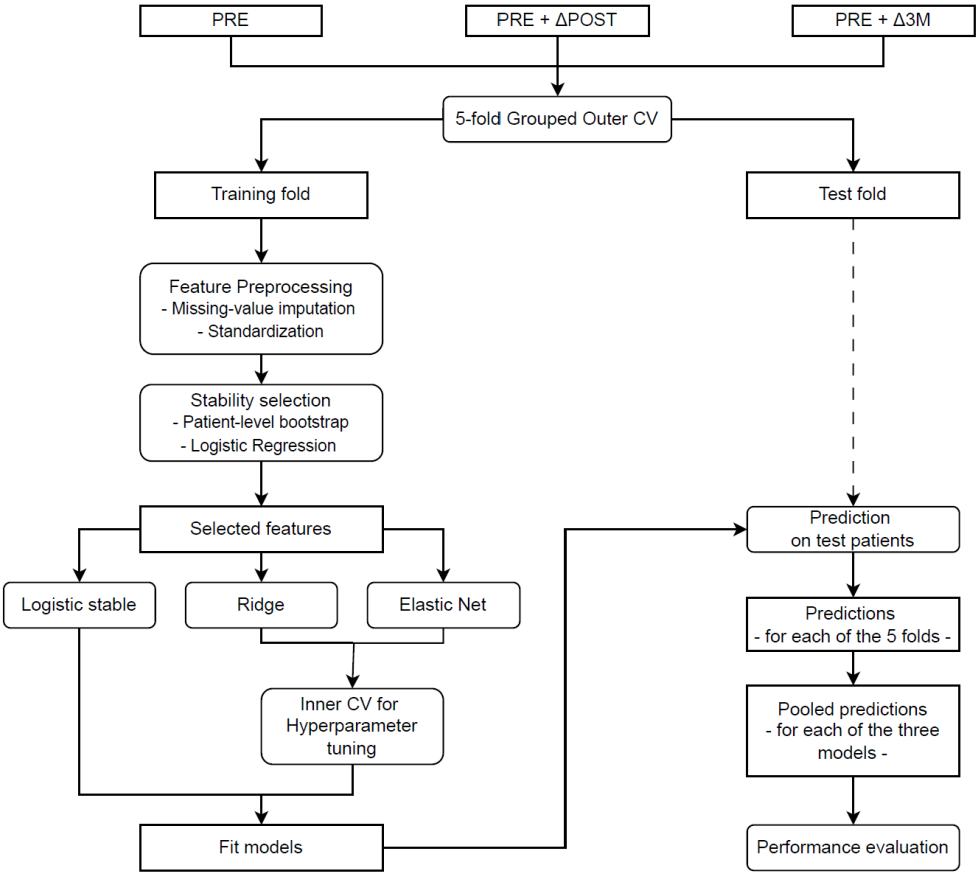


Figure 2.21: Schematic representation of the toxicity-classification workflow, including patient-grouped cross-validation, stability-based feature selection, and evaluation of Logistic Stable, Ridge, and ElasticNet classifiers.

2.4.4 Performance Evaluation Metrics

Classification performance was evaluated according to the binary or multiclass nature of the corresponding analysis.

For longitudinal timepoint classification, confusion matrices and class-specific precision, recall, and F1-scores were used to characterize classification errors across the different acquisition stages. Discriminative performance was further assessed using a one-versus-rest multiclass ROC analysis, reporting class-specific, macro-averaged, and prevalence-weighted ROC-AUC values.

For toxicity discrimination, model evaluation was organized into three complementary levels:

- **Predictive performance.** The primary discrimination metric was the pooled

out-of-fold ROC-AUC, obtained by concatenating the predicted probabilities generated for the held-out patients across the five outer cross-validation folds. In this way, each prediction used for performance assessment was generated by a model that had not been trained on data from the corresponding patient. Accuracy, precision, recall, and F1-score were additionally computed after defining an operating threshold through the Youden index. The Brier score was used to assess the accuracy of the predicted probabilities, while precision–recall curves were examined to account for class imbalance and calibration plots were used to compare predicted probabilities with observed outcome frequencies.

- **Statistical robustness.** To determine whether the observed discrimination could be obtained in the absence of a true association between imaging biomarkers and toxicity, a patient-grouped permutation test was performed. Toxicity labels were randomly reassigned at patient level, preserving the dependency among multiple lesions belonging to the same patient, and the complete modelling procedure was repeated for each permutation. This generated an empirical null distribution of ROC-AUC values, against which the observed pooled out-of-fold ROC-AUC was compared. Statistical evidence against chance-level discrimination was quantified through the corresponding empirical p-value.

- **Incremental contribution of OCT-derived features.** The additional predictive information provided by the OCT-derived morphological biomarkers was evaluated by comparing the original OCTA–radiomics feature set with the extended OCT–OCTA–radiomics representation. The comparison was restricted to lesions available in both datasets and identical patient-grouped cross-validation partitions were used for the two feature sets. The change in discrimination was quantified as

$$\Delta\mathrm{AUC} = \mathrm{AUC}_{\mathrm{OCT+OCTA+RAD}} - \mathrm{AUC}_{\mathrm{OCTA+RAD}}.$$

  Uncertainty in ΔAUC was estimated through paired patient-cluster bootstrap resampling, in which patients were resampled with replacement together with all corresponding lesions. The resulting bootstrap distribution was used to derive the 95% confidence interval of the paired difference in ROC-AUC.

Chapter 3

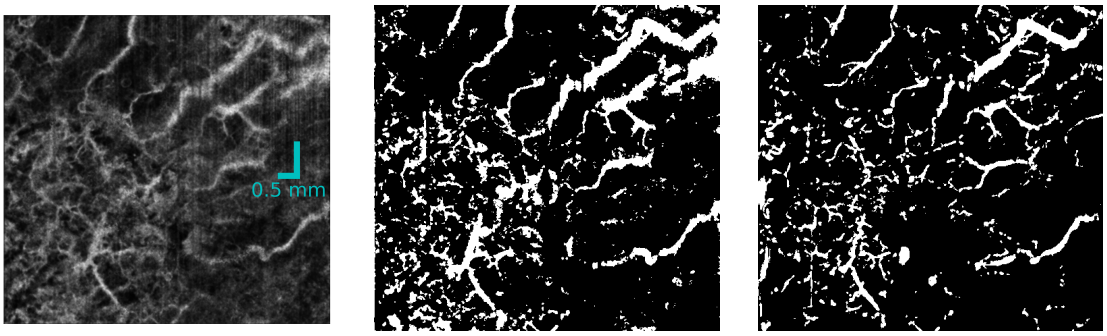
Results

3.1 OCTA Data Optimization

3.1.1 Automatic Segmentation Outcomes

The automatic segmentation framework successfully generated vascular binary masks for the OCTA volumes, enabling the extraction of the corresponding quantitative microvascular features.

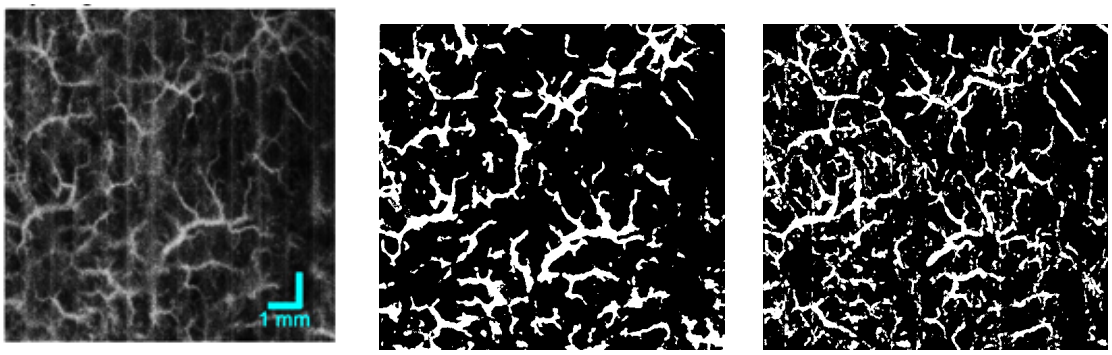
As a first qualitative assessment, the segmentation masks generated by the automatic network were visually compared with the corresponding manual annotations used as reference. Figure 3.1 shows a representative acquisition characterized by good image quality. The automatic segmentation reproduced the main vascular structures identified in the manual reference, with only minor local differences between the two segmentation approaches.



(a) **(b)** **(c)**

Figure 3.1: Representative comparison between manual and automatic vascular segmentation in an OCTA acquisition characterized by good image quality. (a) Original OCTA en-face projection. (b) Manual vascular segmentation used as reference. (c) Corresponding segmentation obtained using the automatic network.

Larger discrepancies between the two segmentation approaches were observed in acquisitions affected by motion artifacts. Figure 3.2 shows a representative example of an artifact-affected acquisition. Compared with the manual reference, the automatic segmentation identified a larger number of small and fragmented structures, particularly within regions characterized by reduced image quality. In some regions, these structures exhibited a predominantly vertical arrangement.



(a) **(b)** **(c)**

Figure 3.2: Representative comparison between manual and automatic vascular segmentation in an OCTA acquisition affected by motion artifacts. (a) Original OCTA en-face projection. (b) Manual vascular segmentation used as reference. (c) Corresponding segmentation obtained using the automatic network.

The quantitative features extracted from the automatically segmented volumes

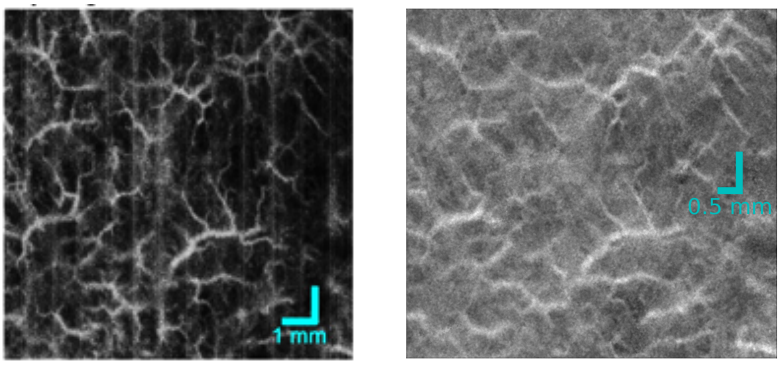
were subsequently evaluated using the longitudinal classification framework described in Section 2.4.1. The original automatic segmentation achieved an overall longitudinal discrimination performance comparable to that obtained from the manual reference. The complete quantitative comparison, including the optimized dataset obtained after targeted motion-artifact correction, is presented in Table 3.1.

3.1.2 Impact of Motion Artifact Correction on Longitudinal Classification

Following the evaluation of the automatic segmentation, acquisitions affected by relevant motion artifacts were identified and subjected to the correction procedure described in Section 2.2. The corrected volumes were then combined with the original uncorrupted acquisitions to generate the optimized dataset.

The number of motion artifacts differed across acquisition timepoints, with the highest proportion observed at 6M (71%).

Motion artifacts were mainly visible in the OCTA en-face projections as hyperreflective vertical stripe-like structures extending across the image. Figure 3.3 shows a representative corrupted acquisition before and after application of the motion-artifact correction procedure. Following correction, the intensity of the stripe-like patterns was reduced and the underlying vascular structures became more clearly visible.



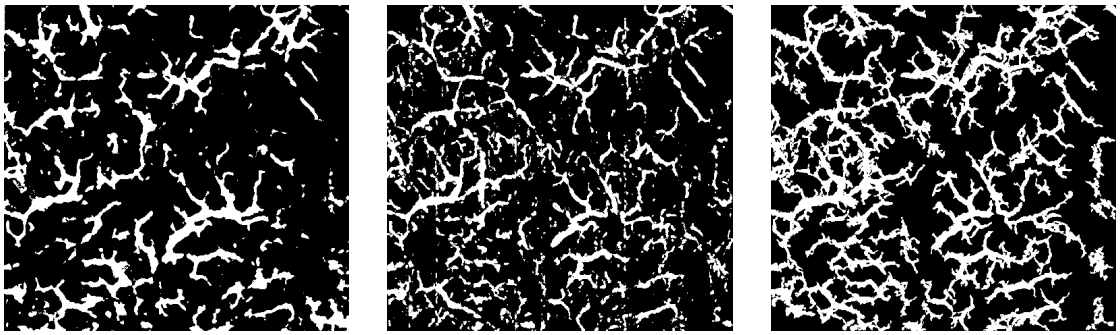
(a) **(b)**

Figure 3.3: Representative OCTA en-face projection affected by motion artifacts before and after correction. (a) Original acquisition showing prominent stripe-like motion artifacts. (b) Corresponding OCTA projection after application of the motion-artifact correction procedure.

The automatic segmentation network was subsequently applied to the corrected volumes. Figure 3.4 compares the manual reference segmentation, the automatic

segmentation obtained from the original motion-corrupted acquisition, and the automatic segmentation obtained after artifact correction for the same representative case.

The overall vascular organization remained comparable among the three segmentation outputs. After artifact correction, the automatic segmentation showed a more continuous representation of several vascular structures that appeared fragmented in the segmentation of the original corrupted volume and included numerous vascular structures that were not retained in the manual reference.



(a) **(b)** **(c)**

Figure 3.4: Representative comparison of vascular segmentation for an OCTA acquisition affected by motion artifacts. (a) Manual reference segmentation. (b) Automatic segmentation of the original motion-corrupted volume. (c) Automatic segmentation obtained after motion-artifact correction.

Although artifact correction generally produced a more continuous vascular representation, heterogeneous segmentation outcomes were observed across the dataset. Figure 3.5 shows a representative case in which a marked discrepancy was found between the manual and automatic segmentation outputs. A large region included in the manual vascular segmentation was not identified by the automatic network, which retained only a limited number of dispersed vascular structures in the same region.

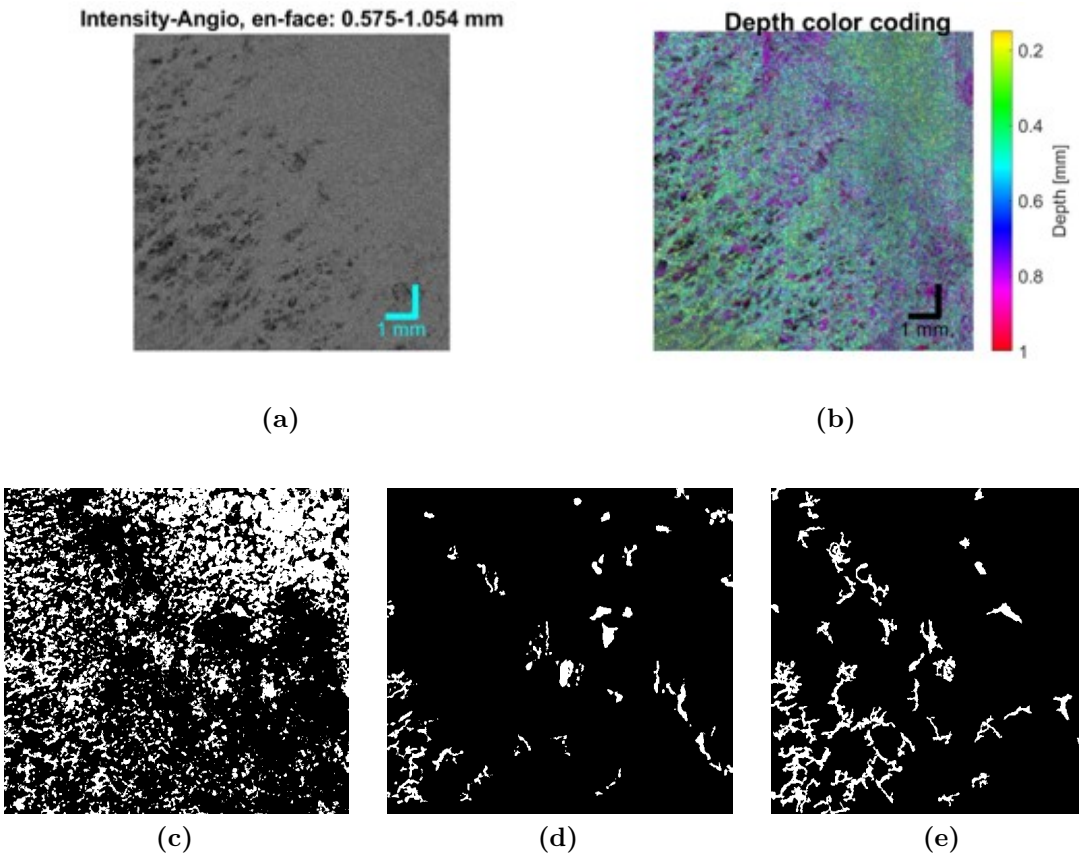


Figure 3.5: Representative challenging OCTA acquisition showing a marked discrepancy between manual and automatic vascular segmentation. (a) Original en-face OCTA projection. (b) Corresponding depth-color-coded representation. (c) Manual reference segmentation. (d) Automatic segmentation of the original acquisition. (e) Automatic segmentation after application of the motion-artifact correction procedure.

The image degradation of this acquisition was not primarily associated with the vertical stripe-like pattern targeted by the motion-artifact correction procedure. In fact, application of the correction algorithm did not produce a substantial change in the corresponding automatic segmentation.

To quantitatively evaluate the effect of the optimization procedure, the same longitudinal classification framework described in Section 2.4.1 was applied to the three feature sets: the one obtained from the manual reference segmentations, the second obtained from the original automatic segmentations, and the last one from the optimized automatic dataset, in which motion-corrupted volumes were replaced by their corrected counterparts.

The macro-averaged AUC was used as the primary metric for the overall multi-class classification performance. In addition, one-vs-rest AUC values were calculated for each acquisition timepoint to evaluate class-specific discrimination. The results obtained for the three processing conditions are summarized in Table 3.1.

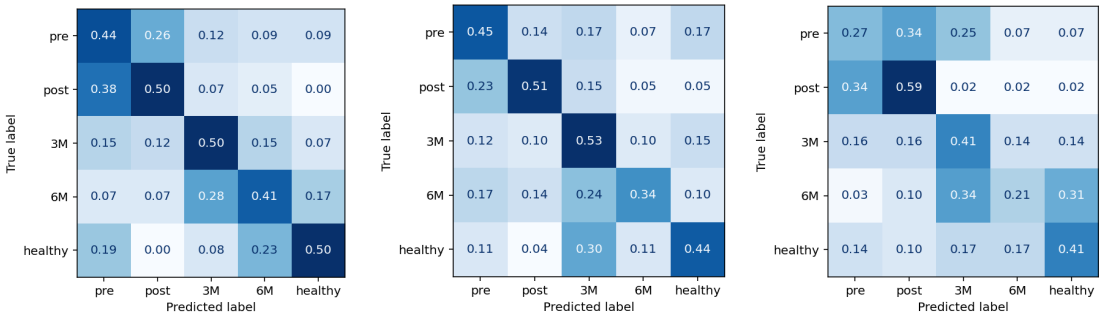
Table 3.1: Longitudinal classification performance obtained using features derived from the manual reference segmentation, the original automatic segmentation, and the optimized automatic segmentation.

Segmentation	**PRE**	**POST**	**3M**	**6M**	**HEALTHY**	**Macro-AUC**
Manual reference	0.67	0.77	0.77	0.79	0.79	0.76
Automatic	0.71	0.77	0.80	0.73	0.80	0.76
Automatic – optimized	0.63	0.82	0.74	0.70	0.84	0.75

Compared with the manual reference, the original automatic pipeline achieved the same macro-AUC (0.76). At class level, the automatic segmentation produced higher AUC values for PRE, 3M and HEALTHY, the same AUC for POST and a lower value for 6M. Following targeted motion-artifact correction, the macro-AUC decreased by 0.01. Relative to the original automatic dataset, the optimized dataset showed higher AUC values for POST and HEALTHY, while lower values were obtained for PRE, 3M, and 6M.

To complement the probability-based discrimination assessed by the AUC, normalized confusion matrices were generated. Unlike the OvR AUC, which is based on the ranking of predicted probability scores for each class against the remaining classes, the confusion matrices reflect the final class assigned to each volume. Consequently, changes in class-specific AUC are not necessarily accompanied by corresponding changes in the diagonal elements of the confusion matrix.

Figure 3.6 shows the normalized confusion matrix obtained for the three features sets. For direct comparison, all three matrices are reported using the same class order and normalization criterion.



(a) **(b)** **(c)**

Figure 3.6: Normalized confusion matrices obtained from the longitudinal classification analysis using features derived from: (a) manual reference segmentations; (b) original automatic segmentations; and (c) optimized automatic segmentations after targeted motion-artifact correction. Each row is normalized by the number of observations belonging to the corresponding true class.

The distribution of the final multiclass predictions differed among the three processing conditions. Following motion-artifact correction, POST showed the largest increase in correct classification rate compared with the manual reference. Conversely, 6M showed the largest decrease, with misclassifications occurring predominantly toward 3M and HEALTHY. Confusion between PRE and POST remained prominent across the three processing conditions, although the direction and magnitude of these misclassifications varied.

The comparison of the three processing conditions therefore provided a quantitative assessment of both the effect of automatic segmentation and the subsequent impact of targeted motion-artifact correction on longitudinal classification performance.

3.2 OCT Morphological Analysis

3.2.1 OCT Segmentation and Morphological Maps

The OCT processing pipeline was applied to the reconstructed volumes to obtain the skin surface and DEJ segmentations required for the extraction of the structural descriptors introduced in this work. Figure 3.7 shows a representative result of the complete processing procedure. The detected skin surface followed the superficial tissue profile, while the DEJ segmentation delineated the deeper interface between the epidermis and dermis. The resulting 3D surfaces were converted into quantitative maps describing superficial skin and DEJ morphology.

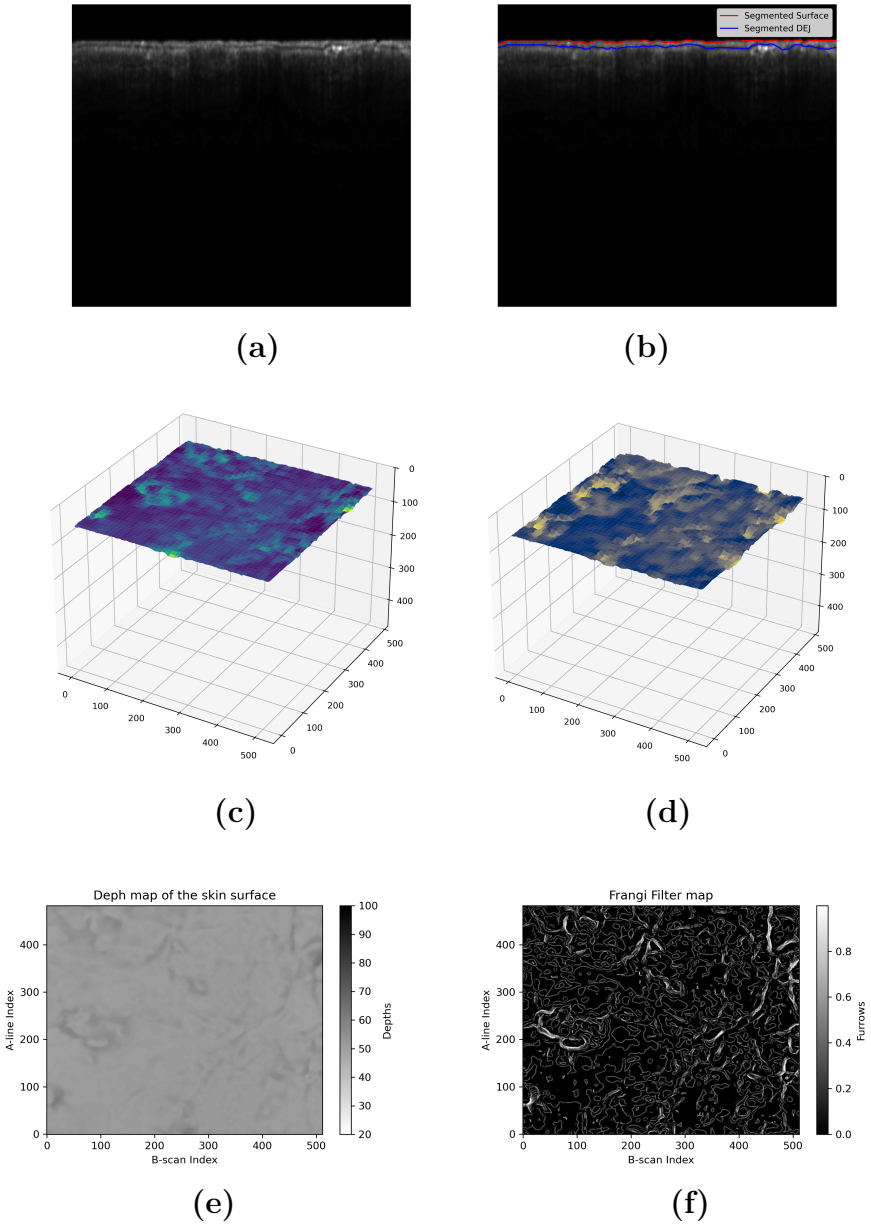


Figure 3.7: Representative outputs of the OCT morphological processing pipeline for the same acquisition. (a) Original OCT B-scan. (b) Corresponding B-scan with the detected skin surface and dermo-epidermal junction (DEJ) superimposed. (c) Three-dimensional representation of the segmented skin surface. (d) Three-dimensional representation of the segmented DEJ surface. (e) Surface-depth map derived from the reconstructed skin morphology. (f) Frangi-derived map used for the characterization of the superficial furrow pattern.

In most volumes, the filtering procedure enhanced the contrast of the hyporeflective layer associated with the DEJ. Nevertheless, poorly defined DEJ contrast remained in some acquisitions, as illustrated in Figure3.8.

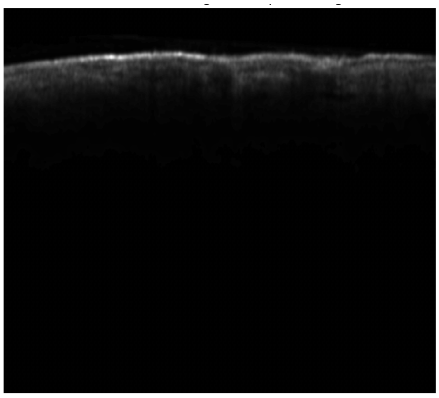
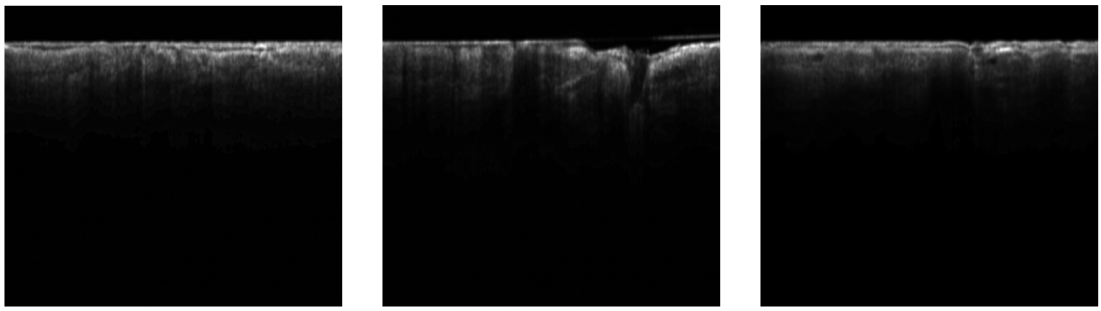


Figure 3.8: Representative example of non visible DEJ in a B-scan

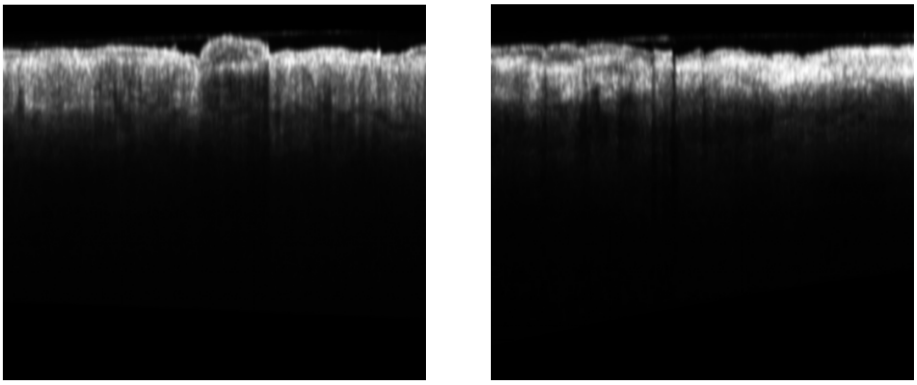
Additional challenging cases exhibited limited areas in which the DEJ appeared interrupted or presented hyporeflective undefined structures.



(a) **(b)** **(c)**

Figure 3.9: Representative OCT B-scans showing showing challenging DEJ visualization. (a) Interruption of the DEJ. (b) Irregular and interrrupted DEJ. (c) Hyporeflective areas below the DEJ.

Superficial crust-like structures were also observed in selected POST and 3M acquisitions. They appeared as heterogeneous, highly scattering and irregular superficial regions that could influence skin-surface reconstruction and the resulting epidermal thickness and furrow-depth measurements. Representative examples are shown in Figure 3.10



(a) **(b)**

Figure 3.10: Representative OCT B-scans showing crusts structure.

In volumes characterized by partial tissue detachment from the optical window, the additional processing branch(* in figure 2.5) was applied. A representative case is shown in Figure 3.11.

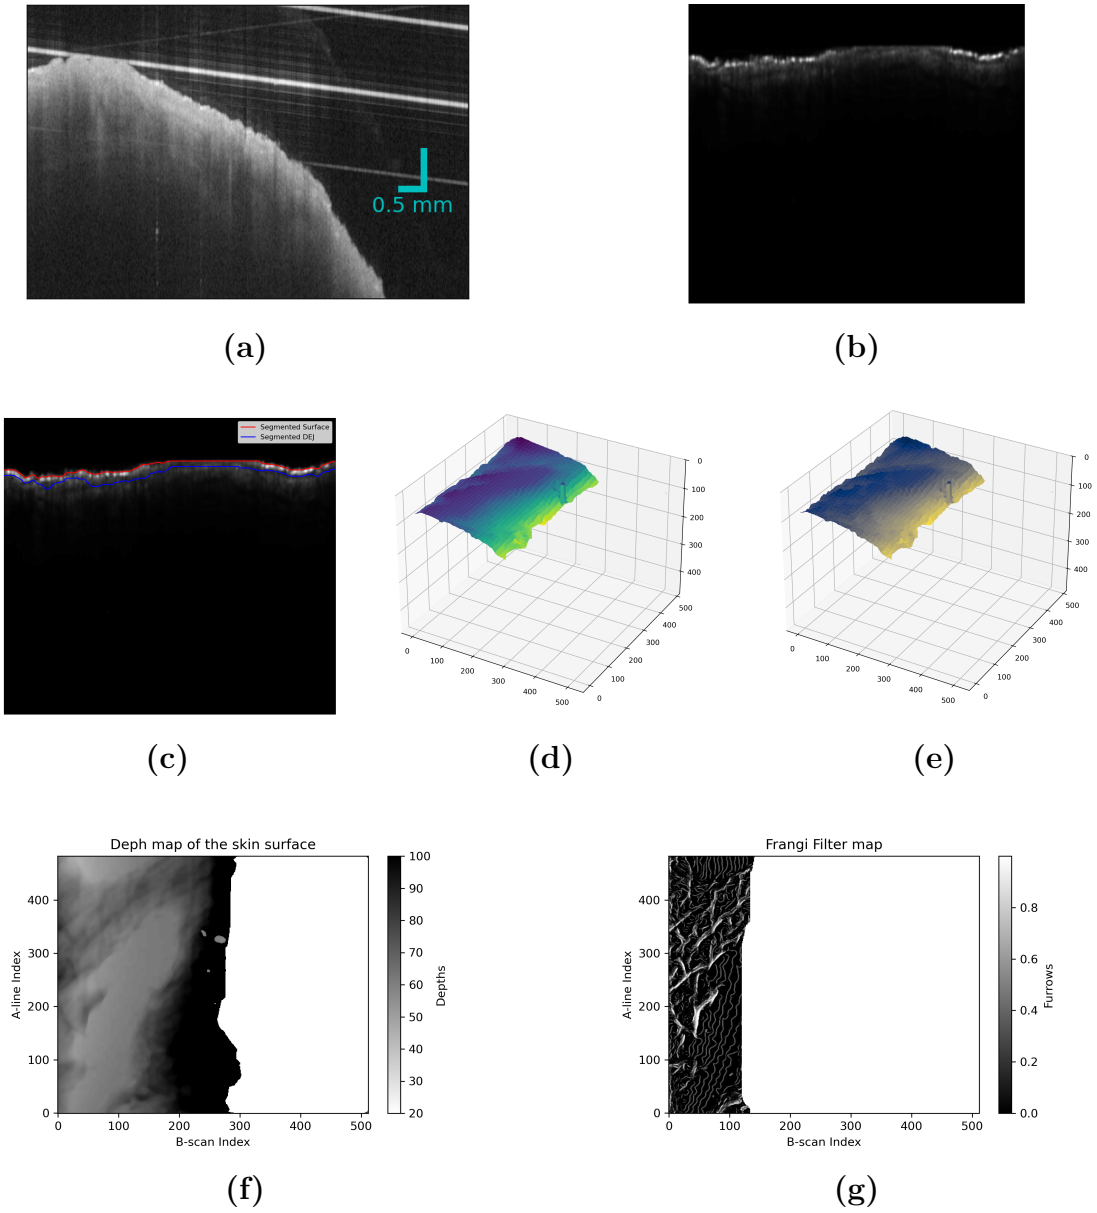
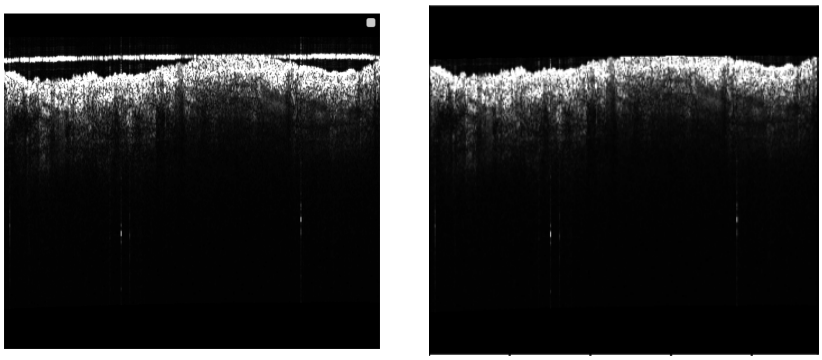


Figure 3.11: Representative outputs of the OCT morphological processing pipeline for the same acquisition. (a) Original OCT cross-scan section. (b) Original B-scan section after glass artifact removal. (c) B-scan with the detected skin surface and dermo-epidermal junction (DEJ) superimposed. (d) Three-dimensional representation of the segmented skin surface after NaN mask application. (e) Three-dimensional representation of the segmented DEJ surface after NaN mask application. (f) Surface-depth map derived from the reconstructed skin morphology. (g) Frangi-derived map used for the characterization of the superficial furrow pattern.

The cross-sectional view highlights the curvature of the tissue beneath the optical window and the presence of multiple reflective glass interfaces, which required the additional artifact-removal procedure. A NaN mask was applied to exclude regions characterized by complete tissue detachment, whereas partially detached regions were retained when sufficient information remained available. Because of these regions, glass-layer removal was applied before the subsequent segmentation steps. Figure 3.12 shows a representative partially detached region before and after glass-artifact removal.



(a) **(b)**

Figure 3.12: Representative OCT B-scan showing partial detachment area with the glass (a) B-scan before glass artifact removal (b) B-scan after glass artifact removal

Regions excluded through the NaN mask were removed from the surface estimation, and the final segmentation was restricted to the valid tissue region.

The resulting segmentations and morphological maps were used to extract the quantitative OCT descriptors subsequently analyzed across acquisition timepoints and toxicity classes.

3.2.2 Distribution of OCT Morphological Features

Initially, the OCT-derived morphological features were evaluated based on their distributions. The distributions were first examined across acquisition timepoints to characterize their longitudinal behavior and identify potential time-dependent differences. Figure 3.13 shows the distributions of mean epidermal thickness, epidermal thickness standard deviation, furrow density, mean furrow depth and maximum furrow depth at PRE, POST, 3M, and 6M, together with the values obtained from healthy skin acquisitions.

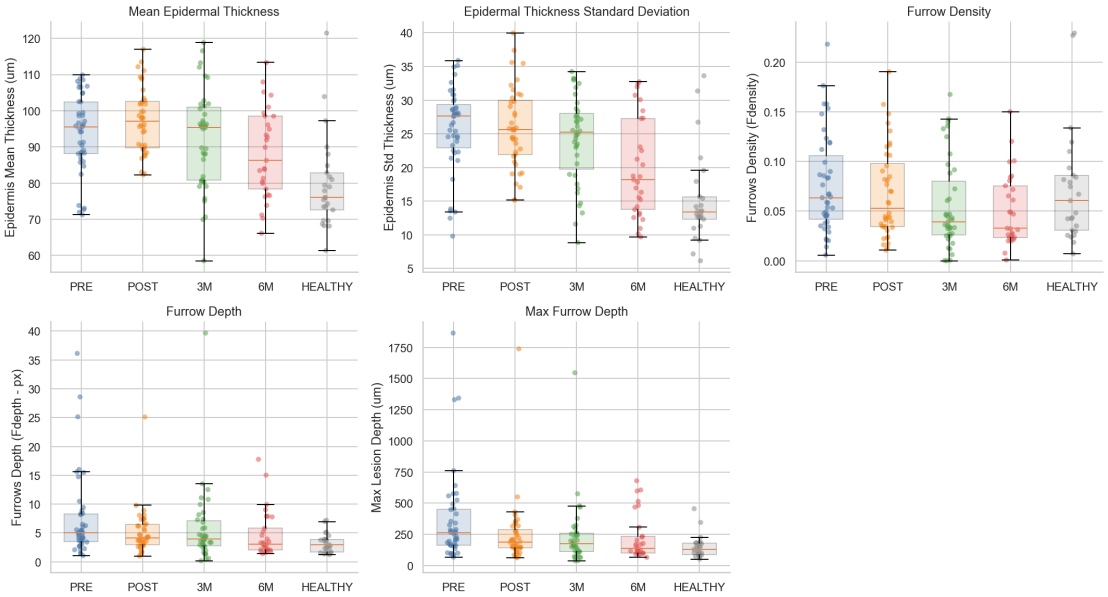


Figure 3.13: Distribution of OCT-derived morphological features at PRE, POST, 3M, 6M, and in healthy skin.

Mean epidermal thickness showed comparable median values at PRE, POST, and 3M, followed by a lower median at 6M. Healthy acquisitions showed the lowest median value among the acquisition conditions. A similar longitudinal pattern was observed for epidermal thickness standard deviation, with lower median values at 6M and in healthy skin compared with the earlier treatment-related timepoints.

Furrow density showed overlapping distributions across the different timepoints, with substantial within-group variability. Furrow depth also showed overlapping distributions, with lower median values at 6M and in healthy acquisitions. Outliers were present at multiple timepoints.

To further examine the behavior of the morphological features and visualize possible separation in distribution, the treated acquisitions were stratified according to toxicity class. Figure 3.14 reports the group-level evolution of the five features for toxicity classes 0 and 1, while the healthy measurements are shown as an independent reference.

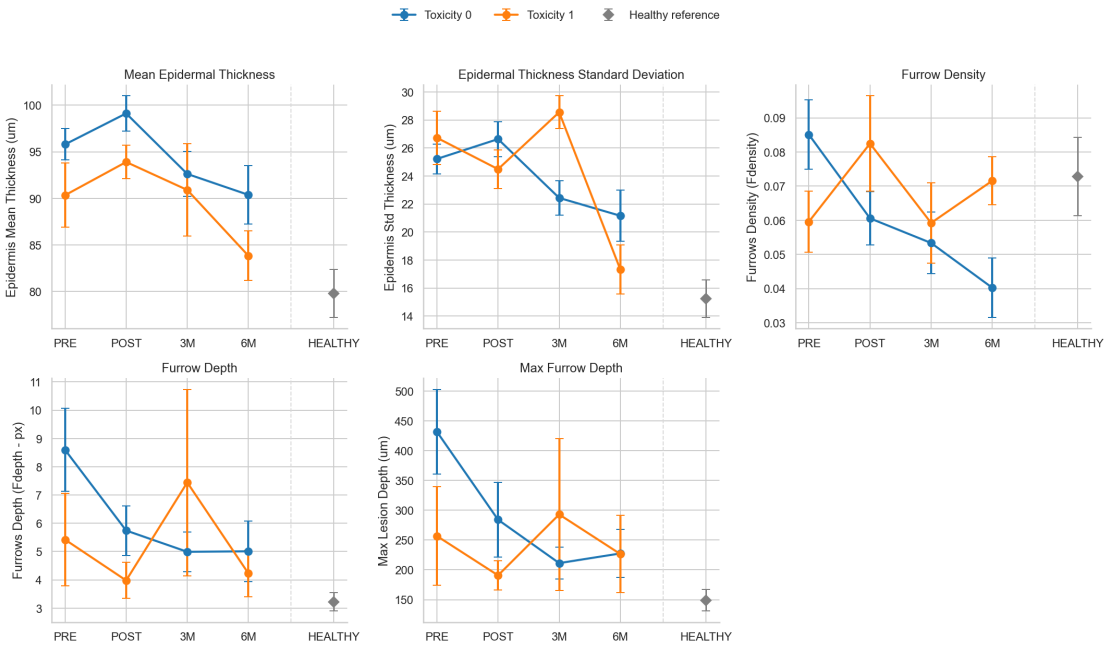


Figure 3.14: Longitudinal evolution of OCT-derived morphological features according to toxicity class. Healthy skin measurements are reported separately as a reference and are not connected to the longitudinal trajectories.

The two toxicity groups exhibited partially different longitudinal trajectories, although substantial overlap remained. Mean epidermal thickness followed a similar temporal trend in both groups, increasing from PRE to POST and subsequently decreasing toward 6M, with consistently lower group-level values for toxicity class 1. Greater between-group differences in temporal behavior were observed for epidermal thickness standard deviation and the furrow-related descriptors, particularly at 3M.

Individual trajectories of the lesions were additionally examined in Figure 3.15, together with the mean trajectory per each group.

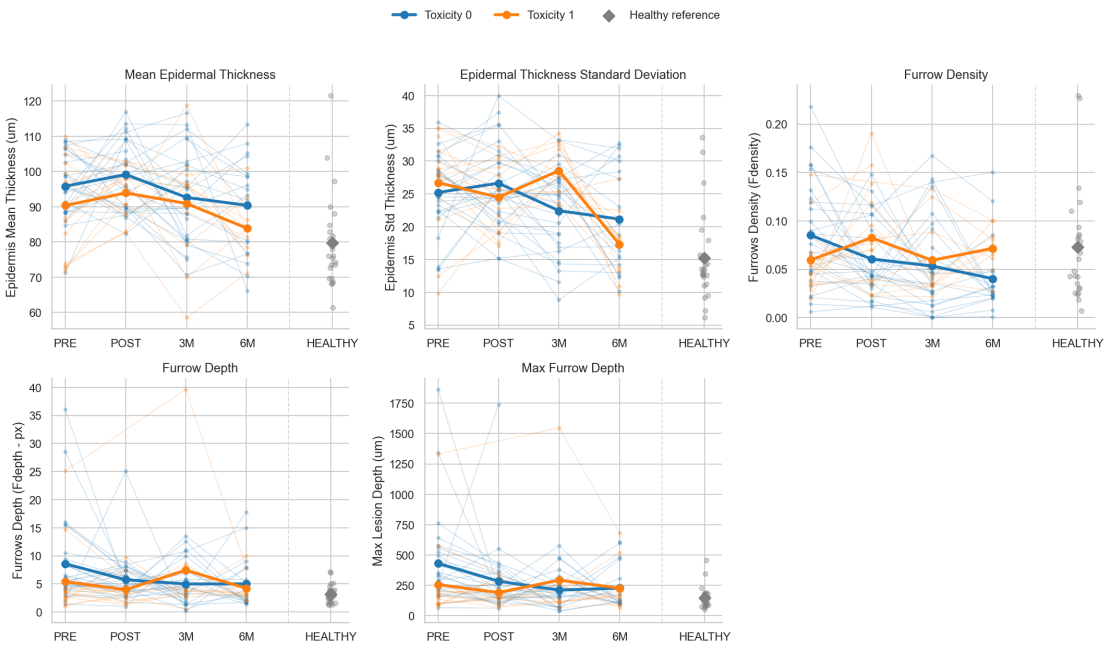


Figure 3.15: Individual and group-level longitudinal evolution of OCT-derived morphological features according to toxicity class. Thin lines represent individual lesion trajectories, while thick lines and markers represent the corresponding group-level values for toxicity classes 0 and 1. Healthy skin measurements are displayed separately as an independent reference.

The individual trajectories revealed substantial inter-lesion variability within both toxicity classes. This variability was particularly evident for furrow depth features, where a limited number of acquisitions showed higher values than the majority of the observations.

Toxicity 0 group showed a more regular behavior, with a decrease from POST to 6M, whereas toxicity 1 exhibited greater variability and did not follow a consistent pattern.

Healthy acquisitions were included as an independent reference and were not assigned a toxicity class, since they corresponded to clinically unaffected skin rather than irradiated lesions.

Overall, substantial overlap between toxicity classes was observed, indicating that the univariate distributions could not provide clear discrimination.

3.3 Contribution of OCT Morphological Features to Quantitative Classification

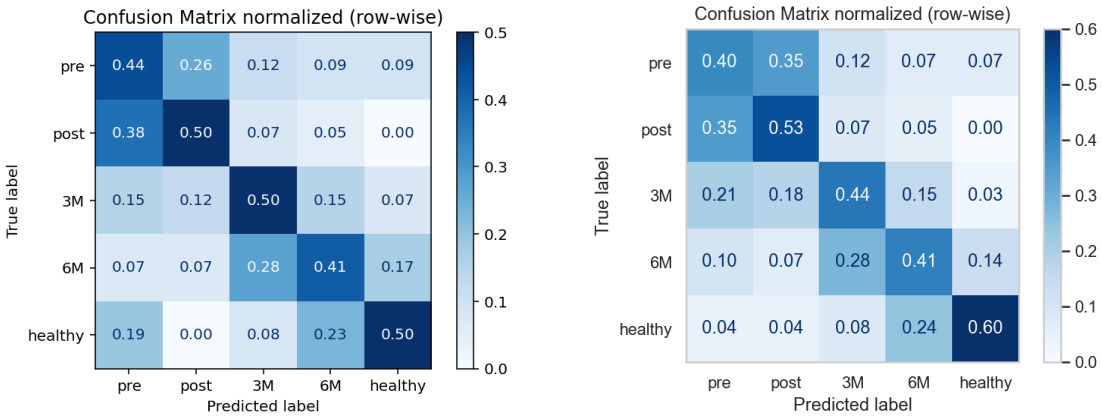
3.3.1 Longitudinal Classification: OCTA/Radiomics vs. OCT/OCTA/Radiomics

To assess the contribution of the OCT-derived morphological features to the characterization of longitudinal treatment-related changes, the classification framework described in Section 2.4.1 was applied to two different feature sets.

The reference feature set included OCTA-derived and radiomic descriptors (OCTA/Radiomics), whereas the extended feature set additionally included the OCT-derived features introduced in the present work (OCT/OCTA/Radiomics). The same classification procedure and target classes were used for both feature sets, allowing their performances to be directly compared.

Figure 3.16 shows the normalized confusion matrices obtained for the two feature representations. Differences in the classification pattern were observed following the introduction of the OCT morphological features.

The most pronounced change following the inclusion of OCT features was observed for the HEALTHY class, which showed an increased correct classification rate. 3M final prediction showed instead the largest decrease in correct classification rate.



(a) **(b)**

Figure 3.16: Normalized confusion matrices obtained from the longitudinal classification analysis using: (a) OCTA-derived microvascular and radiomic features (OCTA/Radiomics); and (b) the extended feature set additionally including OCT-derived morphological features (OCT/OCTA/Radiomics).

Classification performance was further evaluated using the macro-AUC and the one-vs-rest AUC calculated for each acquisition timepoint. The results are reported in Table 3.2.

Table 3.2: Longitudinal classification performance obtained using the OCTA/Radiomics and OCT/OCTA/Radiomics feature sets.

Feature set	PRE	POST	3M	6M	HEALTHY	Macro-AUC
OCTA/Radiomics	0.67	0.77	0.77	0.79	0.79	0.76
OCT/OCTA/Radiomics	0.71	0.79	0.77	0.80	0.82	0.78

The inclusion of OCT-derived morphological features resulted in a change in macro-AUC from 0.76 for the OCTA/Radiomics feature set to 0.78 for the extended OCT/OCTA/Radiomics feature set. At the class level, the largest improvement was observed for HEALTHY timepoint, whereas 3M showed no variations following the addition of the OCT-derived features.

To investigate the contribution of individual variables to the classification, the feature ranking generated within the classification framework was also examined. Figure 3.17 shows the highest-ranking features obtained for the extended OCT/OCTA/Radiomics feature set.

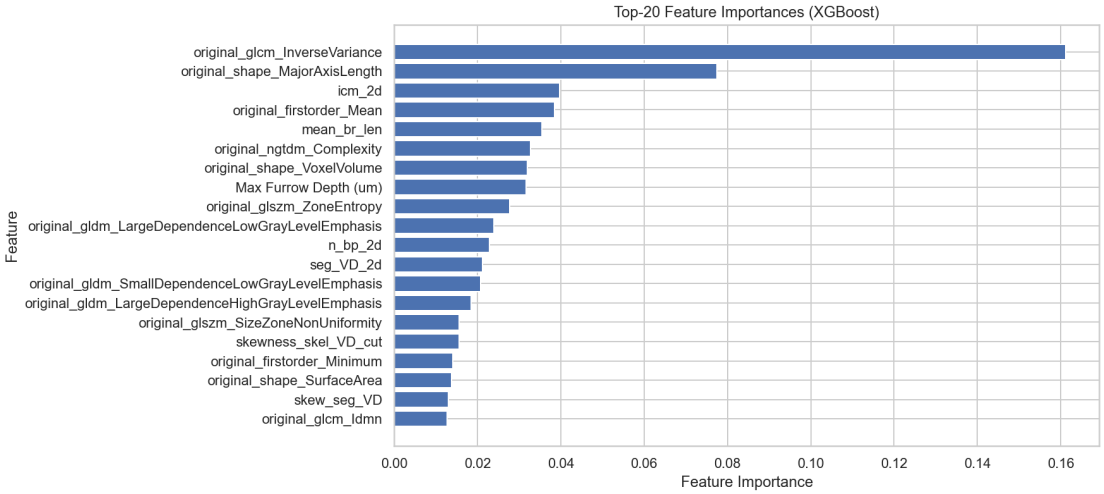


Figure 3.17: Ranking of the highest-contributing features in the longitudinal classification model using the extended OCT/OCTA/Radiomics feature set.

The inclusion of the OCT-derived features also modified the overall feature-ranking profile. Max Furrow depth was retained among the highest-ranking variables of the extended model, together with OCTA-derived and radiomic descriptors.

3.3.2 Toxicity Classification and Contribution of OCT Morphological Features

The toxicity classification analysis was performed to investigate whether quantitative imaging features could discriminate between the two toxicity classes defined in Section 2.4.2 and to assess whether the addition of OCT-derived morphological information improved classification performance.

To investigate the discriminative information available at different stages of treatment, three temporal feature representations were considered: baseline features acquired before treatment (PRE), changes between PRE and immediately post-treatment (PRE + ΔPOST), and changes between PRE and the 3-month follow-up (PRE + Δ3M).

For each temporal representation, feature selection and model training were performed according to the classification framework described in Section 2.4.3. Three regularized linear classification models were evaluated: stable logistic regression, ridge logistic regression, and elastic-net logistic regression. Model performance was assessed using ROC-AUC, accuracy, precision, recall, F1-score, and Brier score.

First, the extended OCT/OCTA/Radiomics feature set was evaluated.

The classification results obtained are summarized in Table 3.3.

Table 3.3: Toxicity classification performance obtained using the extended OCT/OCTA/Radiomics feature set across the three temporal feature representations and classification models. The highest ROC-AUC obtained within the extended feature set is highlighted.

Dataset	Model	ROC-AUC	Accuracy	Precision	Recall	F1-score	Brier score
PRE	Logistic Stable	0.685	0.627	0.481	0.866	0.619	0.244
PRE	Ridge	0.657	0.744	0.666	0.533	0.592	0.254
PRE	Elastic Net	0.707	0.720	0.615	0.533	0.571	0.278
PRE + ΔPOST	Logistic Stable	0.665	0.717	0.588	0.714	0.645	0.257
PRE + ΔPOST	Ridge	0.737	0.717	0.588	0.714	0.645	0.234
PRE + ΔPOST	Elastic Net	0.720	0.743	0.625	0.714	0.666	0.254
PRE + Δ3M	Logistic Stable	0.877	0.810	0.642	0.818	0.720	0.134
PRE + Δ3M	Ridge	**0.881**	0.864	0.750	0.818	0.782	0.205
PRE + Δ3M	Elastic Net	0.655	0.540	0.392	1.000	0.564	0.199

Among the configurations evaluated using the extended feature set, the highest ROC-AUC was obtained for the PRE + Δ3M representation using the ridge classifier, which achieved a ROC-AUC of 0.88.

The statistical significance of the observed classification performance was further evaluated using the permutation procedure.

Figure 3.18 shows the resulting null distribution together with the ROC-AUC obtained using the original toxicity labels. The observed ROC-AUC of 0.88 was higher than all the values obtained after label permutation, corresponding to a

permutation p-value of 0.0099.

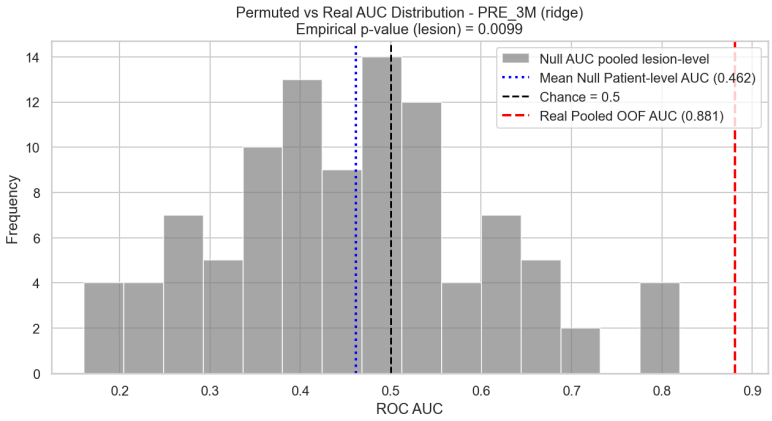


Figure 3.18: Permutation analysis of the PRE + Δ3M ridge toxicity classifier using the extended OCT/OCTA/Radiomics feature set. The distribution represents the ROC-AUC values obtained after random permutation of the toxicity labels, while the vertical line indicates the ROC-AUC obtained using the original labels (ROC-AUC = 0.88; permutation $p = 0.0099$).

Following the evaluation of the extended feature set, a comparative analysis was performed to assess the contribution of the newly introduced OCT-derived morphological features. For this purpose, the classification performance obtained using the OCT/OCTA/Radiomics feature set was compared with that obtained using the reference feature set (OCTA/Radiomics).

The same temporal representations and classification framework were used for both feature sets, allowing their discriminative performances to be directly compared. Table 3.4 summarizes the results obtained for PRE, PRE + ΔPOST, and PRE + Δ3M.

Table 3.4: Comparison of toxicity classification performance between the OCTA/Radiomics and OCT/OCTA/Radiomics feature sets. The PRE dataset included 43 lesions from 19 patients, the PRE + ΔPOST dataset included 39 lesions from 17 patients, and the PRE + Δ3M dataset included 37 lesions from 17 patients. Differences in ROC-AUC (ΔAUC), 95% confidence intervals (CI), and p-values were obtained from the paired comparison between the two feature sets.

Dataset	**Model**	**OCTA/ Radiomics**	**OCT/OCTA/ Radiomics**	**ΔAUC**	**95% CI**	**p-value**
PRE	Logistic Stable	0.633	0.685	0.052	[-0.060; 0.160]	0.339
PRE	Ridge	0.600	0.657	0.057	[-0.137; 0.258]	0.563
PRE	Elastic Net	0.635	0.707	0.071	[-0.114; 0.272]	0.412
PRE + ΔPOST	Logistic Stable	0.631	0.665	0.034	[-0.074; 0.186]	0.626
PRE + ΔPOST	Ridge	0.622	0.737	0.114	[-0.070; 0.377]	0.254
PRE + ΔPOST	Elastic Net	0.608	0.720	0.111	[-0.114; 0.374]	0.370
PRE + Δ3M	Logistic Stable	0.881	0.877	-0.003	[-0.141; 0.087]	0.938
PRE + Δ3M	Ridge	0.905	0.881	-0.024	[-0.181; 0.062]	0.628
PRE + Δ3M	Elastic Net	0.596	0.655	0.059	[0.000; 0.154]	0.672

The addition of OCT-derived morphological features produced different changes in ROC-AUC depending on the temporal representation and classification model. For the PRE dataset, the ROC-AUC increased for all three classifiers. Similarly, for the PRE + ΔPOST representation, positive ΔAUC values were observed for all three classifiers, with the largest increase obtained using Ridge (ΔAUC = 0.114).

For the PRE + Δ3M representation, smaller differences were observed between the two feature sets. Logistic Stable produced almost the same ROC-AUC for both feature representations, while Elastic Net showed an increase. Ridge, which achieved the highest ROC-AUC among the evaluated configurations, showed a slight decrease from 0.90 with OCTA/Radiomics to 0.88 with OCT/OCTA/Radiomics (ΔAUC = -0.02).

The 95% confidence intervals of ΔAUC included zero for all comparisons, and none of the paired comparisons reached statistical significance ($p > 0.05$).

To quantify the uncertainty associated with the differences observed between the two feature representations, a paired bootstrap analysis was performed. Since the predictions generated by the OCTA/Radiomics and OCT/OCTA/Radiomics models were obtained from the same observations, paired resampling preserved the correspondence between predictions and allowed the difference in ROC-AUC between the two feature sets to be directly estimated.

The visualization of resulting paired bootstraps is displayed in Figure 3.19

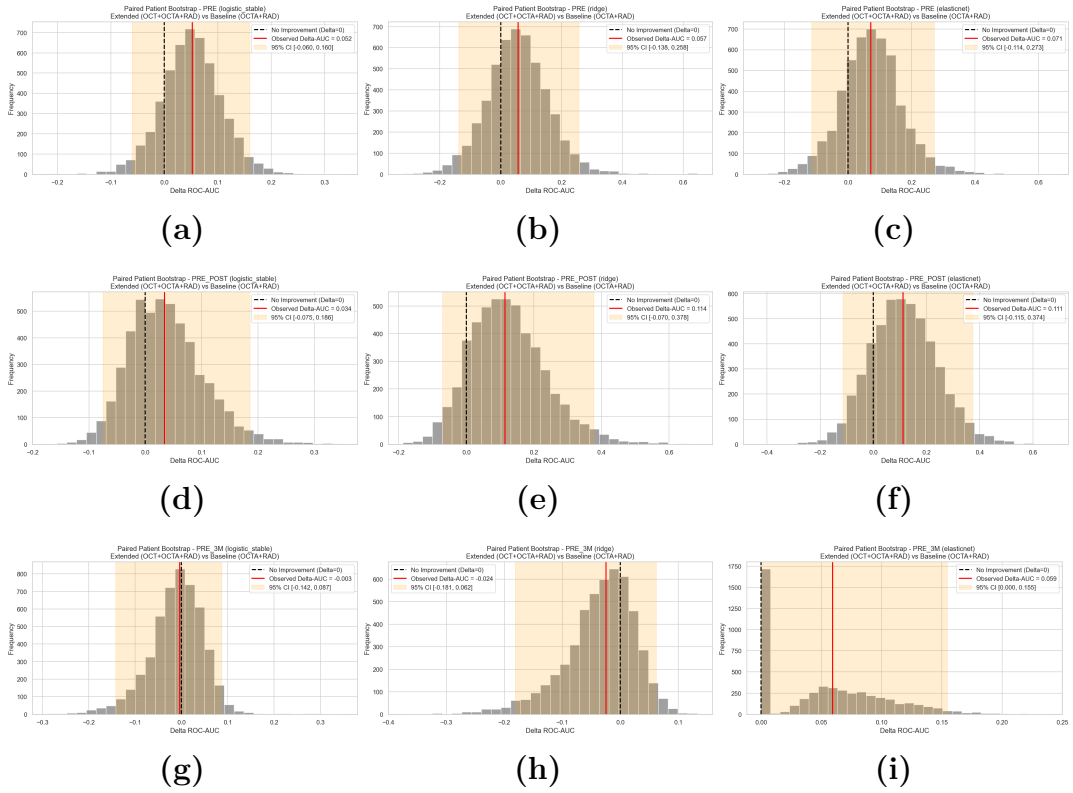


Figure 3.19: Paired bootstrap distributions of the differences in ROC-AUC between the OCT/OCTA/Radiomics and OCTA/Radiomics feature sets. Panels (a–c) correspond to PRE, panels (d–f) to PRE + ΔPOST, and panels (g–i) to PRE + Δ3M. Within each row, Logistic Stable, Ridge, and Elastic Net classifiers are shown from left to right. Positive ΔAUC values indicate higher ROC-AUC after inclusion of OCT-derived morphological features, whereas negative values indicate higher ROC-AUC for the OCTA/Radiomics feature set.

For the PRE + ΔPOST (ridge) model, which showed the largest increase after the introduction of OCT features, the stability selection procedure identified a subset of recurrently retained features, presented in Table 3.5.

Table 3.5: Features selected by the stability-selection procedure for the PRE + ΔPOST OCT/OCTA/Radiomics feature set. Features are reported in decreasing order of selection frequency.

Selected feature	**Selection frequency (%)**
ΔPOST Epidermal Thickness Standard Deviation (μm)	60.0
GLCM Difference Variance	55.0
Shape Sphericity	50.0
ΔPOST Furrows Density	49.0
Epidermis Mean Thickness (μm)	48.0
ΔPOST MR 3D	46.0
GLCM Cluster Shade	44.0
ΔPOST GLCM Difference Variance	40.0
MR 2D	38.0
First-order Skewness	37.0

Notably, three OCT-derived morphological descriptors were retained among the ten most frequently selected features: ΔPOST Epidermal Thickness Standard Deviation was selected with a frequency of 60%, ΔPOST Furrows Density with a frequency of 49.0% and Epidermis Mean Thickness with a frequency of 48%.

Chapter 4

Discussion and Conclusions

4.1 Data Optimization

The introduction of automatic vascular segmentation represents an important methodological step toward increasing the reproducibility, objectivity, and scalability of the OCTA analysis workflow. Qualitative comparison with the manual reference showed that the automatic network was generally able to reproduce the overall vascular organization, particularly for larger and well-defined vessels, whose main branches were consistently identified. The largest discrepancies were observed for smaller vascular structures, which were more frequently retained by the automatic segmentation.

This behavior may partly reflect the intrinsic differences between manual and automated interpretation of complex vascular patterns. The segmentation network was trained to identify recurring morphological characteristics associated with vascular structures, whereas the vascular architecture observed in the clinical dataset can be substantially altered by the presence of the tumor and by treatment-induced remodeling. Consequently, atypical, fragmented, or poorly defined vascular patterns may be interpreted differently by the network and by the human operator. On the other hand, manual segmentation should not be considered an entirely objective ground truth. Although it represents the reference approach adopted in the original study and benefits from the operator's ability to distinguish vascular structures from image noise, it remains inherently dependent on visual interpretation. Therefore, it cannot be excluded that some of the small structures retained by the automatic network may correspond to actual vascular information that was not included in the manual segmentation.

Despite these local differences, the quantitative analysis showed that the overall discriminative information extracted from the automatic segmentations was largely

preserved. The manual and automatic feature sets achieved the same macro-averaged AUC in the longitudinal classification, although differences were observed for individual acquisition timepoints.

In particular, HEALTHY acquisitions were among the best-discriminated classes using automatically derived features. Healthy skin is expected to exhibit a more regular vascular organization than tumor tissue undergoing radiotherapy, potentially resulting in patterns that are more consistently recognisable by the automatic segmentation. Differences in sample size and acquisition quality among timepoints may also have contributed to the observed class-specific behavior.

The subsequent application of motion-artifact correction produced an apparently contrasting result. From a qualitative perspective, the procedure effectively reduced the characteristic vertical stripe-like patterns affecting several OCTA acquisitions. Removal of these artifacts generally improved the visibility and continuity of the underlying vascular structures and reduced the presence of isolated signals in the corresponding automatic segmentation. Nevertheless, this qualitative improvement was not associated to an increase in overall longitudinal classification performance. The macro-averaged AUC slightly decreased after optimization, with heterogeneous class-specific effects: discrimination improved for POST and HEALTHY acquisitions, whereas lower AUC values were obtained for PRE, 3M, and particularly 6M. Several factors may explain this discrepancy.

The correction algorithm was specifically designed to suppress the vertical stripe-like pattern associated with motion artifacts, whereas clinical OCTA acquisitions may also be affected by other acquisition-related distortions. For example, the phenomenon shown in Figure 3.5 could be potentially associated with very small and fragmented capillary structures, which may represent a plausible post-treatment vascular pattern. Their limited dimensions may approach the effective resolution limits of OCTA, reducing their visibility and contributing to incomplete vascular representation. Consequently, correction of the targeted artifact cannot recover information degraded by different mechanisms. Moreover, in severely corrupted regions, suppression of the stripe-like pattern may improve image appearance without restoring the underlying vascular signal required for reliable segmentation and feature extraction. The larger change observed for the 6M class should also be considered in relation to the distribution of motion artifacts across acquisition timepoints. Approximately 71% of the available 6M acquisitions were affected by motion artifacts, representing the highest proportion among the treatment-related timepoints. Consequently, a relatively large fraction of this class was modified by the optimization procedure, potentially contributing to the greater change observed in its downstream classification performance.

In general, motion-artifact correction proved useful for recovering selected corrupted acquisitions, but did not systematically improve classification performance. Within the proposed framework, it should therefore be considered primarily as

a targeted quality-control and data-recovery strategy rather than as a procedure expected to consistently enhance downstream classification.

4.2 Interpretation of OCT/OCTA Quantitative Findings

The OCT-derived morphological features introduced in this study provide quantitative information on structural skin characteristics that complements the information obtained from OCTA and radiomics analysis. Among these parameters, epidermal thickness represents a particularly relevant morphological descriptor, although its interpretation requires consideration of the substantial physiological variability of the epidermis. Previous studies have demonstrated that epidermal thickness depends strongly on the anatomical region under investigation and can also be influenced by individual characteristics such as age and sex [13] [11]. In particular, sun-exposed areas show greater epidermal and stratum corneum thickness than sun-protected areas, while increasing age has been associated with a substantial decrease in epidermal thickness. A systematic review and meta-analysis provided reference estimates according to body site, age, sex and measurement technique [10].

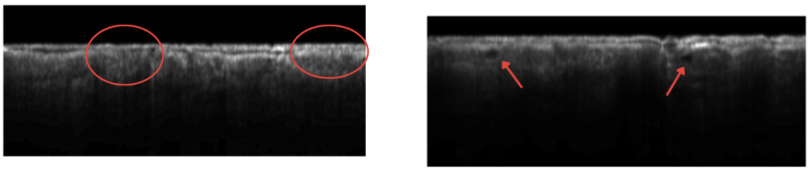
These considerations are relevant for the present dataset, since most of the investigated lesions were located in sun-exposed regions and the cohort predominantly included older patients. Nevertheless, inter-individual variability should be taken into account when evaluating the results.

The epidermal thickness values obtained in this study were overall comparable with the ranges previously reported using OCT, two-photon microscopy, and histological measurements [13] [10]. In contrast to these reference studies, the present cohort consisted of patients affected by NMSC or precancerous lesions and included measurements acquired before and after radiotherapy. Consequently, the measured epidermal thickness may reflect not only physiological variability, but also lesion-related alterations, the presence of superficial crusts or open wound, and structural changes induced by irradiation.

Particular attention should also be given to the reliability of DEJ localization, since this interface defines the lower boundary used for epidermal thickness estimation. Its detection can become challenging when the normal optical contrast between epidermal and dermal tissue is reduced. Previous OCT studies have described variability in the appearance and definition of skin structures, including the DEJ [26]. This issue becomes especially relevant in irradiated skin, where the visibility of the DEJ has been reported to depend on the scattering characteristics of the epidermal and dermal layers and to vary both among patients and during the course of radiotherapy [27]. Inflammatory processes [17] and the presence of

superficial crusts or open wounds can further reduce the optical contrast between epidermis and dermis. In the present dataset, similar difficulties were observed in selected acquisitions, where the DEJ appeared locally poorly demarcated, irregular, or interrupted.

An example of such cases is highlighted in Figure 4.1(a). The irregular apperarance of the DEJ may contribute to uncertainty in epidermal thickness estimation.



(a) **(b)**

Figure 4.1: Representative OCT B-scans highlighting morphological alterations discussed in the text. Red circles indicate regions in which the dermo-epidermal junction (DEJ) appears locally interrupted or poorly defined, while red arrows indicate hyporeflective dermal areas.

The dark areas identified in the dermis in Figure 4.1(b) could be related to the intrinsic structure of BCC pre-radiotherapy, which is characterized by the presence of cysts and cell nests, based on the type (superficial, nodular, micronodular, infiltrative, and mixed variants). These could be associated with hyporeflective regions in OCT images [28] [17].

Radiotherapy-induced alterations of the superficial epidermal layers may represent an additional source of morphological variability. After radiotherapy, phenomena like development of epidermal disordering and thickening have been registered [27].

Surface morphology provided an additional level of structural information. The furrow-related descriptors implemented in this study were based on the three-dimensional skin-surface analysis and Frangi-based identification of furrow-like structures previously proposed for OCT skin roughness assessment [29] [30]. The magnitude of the furrow-depth measurements obtained in the present dataset was approximately comparable with values previously reported using similar OCT-based approaches [29]. However, the distributions included several higher values and outliers, which may partly reflect the fact that previous methodological studies were predominantly performed on healthy skin, whereas the present dataset included tumor-bearing and irradiated tissue characterized by a less regular surface architecture.

Furrow density showed lower values than those previously reported. This difference should not necessarily be interpreted as a purely biological effect. In

the present implementation, the Frangi-filter response was thresholded following visual assessment of the resulting maps in order to retain structures considered representative of actual skin furrows while reducing the contribution of weak or non-specific responses. Consequently, the adopted procedure provided a more restrictive definition of the surface pixels classified as furrows.

The longitudinal analysis showed substantial overlap among acquisition timepoints for most OCT-derived morphological features, indicating that no individual descriptor exhibited a clearly defined temporal pattern. Nevertheless, epidermal thickness, its variability, and furrow-depth descriptors showed a common tendency toward lower values at 6M and in healthy acquisitions compared with the earlier treatment-related timepoints, whereas furrow density remained largely overlapping across acquisition conditions. The shift of several morphological parameters toward the values observed in healthy skin at 6M may be compatible with progressive structural remodeling following treatment, although it should not be interpreted as evidence of complete tissue recovery.

Conversely, the greater variability observed at earlier follow-up stages may reflect the heterogeneous structural response to radiotherapy, consistently with previous OCT observations of radiation-induced epidermal thickening and disruption of normal epidermal organization [27].

Radiation-induced skin toxicity does not develop uniformly, and in the present cohort toxicity was clinically recorded at different stages of follow-up (0, 1, 3 months). Consequently, acquisitions obtained at POST and particularly at 3M may represent heterogeneous stages of tissue response, ranging from limited residual alterations to persistent or more pronounced toxicity-related changes. This temporal heterogeneity provides a plausible explanation for the broader distributions observed at intermediate follow-up stages.

Stratification according to toxicity class did not reveal a clear univariate separation, as substantial overlap and inter-lesion variability were observed for all morphological descriptors. However, lesions classified as toxicity 0 generally showed a more regular evolution toward 6M, whereas toxicity 1 lesions exhibited greater variability and less homogeneous trajectories. Therefore, the dispersion observed at 3M may partly reflect differences in individual tissue response to irradiation.

This would suggest that the relevance of OCT-derived parameters may arise from their complementary contribution within a multivariate framework, as further supported by the subsequent feature-selection and classification analyses.

4.3 Multimodal Imaging and Radiation-Induced Toxicity Assessment

The longitudinal classification analysis provided an initial assessment of the contribution of OCT-derived morphological features to the multimodal framework based on OCTA and radiomic descriptors. Their inclusion resulted in a modest improvement in overall discrimination, with the macro-AUC increasing from 0.76 to 0.78. Although limited at the global level, the effect was more evident for specific acquisition conditions, particularly HEALTHY, for which both the discriminative performance and the correct classification rate increased. These findings suggest that OCT-derived morphology provides information that is at least partially complementary to the vascular and radiomic characterization.

A stronger effect emerged when the temporal evolution of the tissue was considered for toxicity assessment. Baseline features alone provided only moderate discrimination between toxicity classes, whereas the inclusion of longitudinal changes, particularly those occurring between PRE and 3M, substantially increased the discriminative performance. Within the extended OCT/OCTA/Radiomics feature set, the highest performance was obtained with the Ridge model using the PRE + Δ3M representation. This suggests that the response of the tissue to irradiation carries relevant information associated with toxicity that is not fully represented by its baseline characteristics. The permutation analysis performed on the PRE + Δ3M extended model confirmed that the observed discrimination was unlikely to result from a random association between imaging features and toxicity labels.

After having obtained the performances on the final feature set, the incremental contribution of OCT features has been investigated by comparing the performances to those obtained with the OCTA/radiomics feature set. The addition of morphological descriptors improved performance in most model configurations, with the largest increases observed when PRE and particularly PRE + ΔPOST information were considered. The stability-selection analysis of the PRE + ΔPOST Ridge configuration further supports this interpretation. Three OCT-derived descriptors were retained among the most frequently selected variables, with ΔPOST Epidermal Thickness Standard Deviation showing the highest selection frequency overall.

However, the paired bootstrap analysis did not demonstrate statistically significant differences between the two feature sets. Therefore, the present results support a potential complementary contribution of OCT morphology, but do not provide sufficient evidence to establish a significant improvement over OCTA and radiomics alone.

The incremental contribution of OCT became smaller for PRE + Δ3M, despite this representation providing the strongest overall toxicity discrimination. At 3M, vascular and radiomic changes already showed strong discriminative information,

potentially leaving less independent information to be added by the limited set of morphological OCT descriptors. Conversely, when OCTA and radiomic information alone was less discriminative, particularly at earlier stages, the addition of structural information produced larger improvements. The relevance of the 3M timepoint should nevertheless be interpreted cautiously. As shown by the morphological analysis, this stage was characterized by substantial inter-lesion variability, potentially reflecting differences in the timing and evolution of individual tissue responses. Moreover, the PRE + Δ3M analysis included fewer patients and lesions than the baseline analysis, increasing the uncertainty associated with model evaluation. The high discrimination observed at this timepoint therefore represents a promising finding, but requires confirmation in a larger cohort. This analysis should be interpreted as a toxicity discrimination study rather than a predictive model, since toxicity could already have developed by the follow-up timepoints considered.

Overall, the results suggests that longitudinal changes provide more informative characterization of radiation-induced toxicity than baseline features alone. Within this framework, OCT-derived morphology appears to provide its greatest incremental contribution at earlier stages, whereas at 3M much of the discriminative information may already be captured by vascular and radiomic descriptors.

4.4 Limitations and Future Developments

The findings of this study should be interpreted considering several limitations. First, the relatively limited cohort size represents one of the main constraints. The final dataset included 19 patients, with multiple lesions acquired from some subjects. Although the classification frameworks were designed to account for the patient-level structure of the data and avoid information leakage between training and testing sets, the lack of complete independence among lesion-level observations remains a limitation.

Furthermore, the number of available observations decreased when longitudinal differences were considered because because complete longitudinal acquisitions were not available for all patients, further limiting the statistical power of the toxicity analyses. Expanding the cohort would therefore be essential to obtain more precise performance estimates, reduce uncertainty, and determine whether the observed improvements associated with OCT integration can be confirmed statistically. In particular, the absence of significant differences in the paired bootstrap comparison should not be interpreted as evidence that OCT provides no additional information, but rather as an indication that its incremental contribution could not be demonstrated in the present limited cohort.

A second limitation concerns the heterogeneity of image acquisitions. As discussed in the previous sections, the rigid glass interface could not provide uniform contact on all anatomical regions, particularly on curved or irregular surfaces. This resulted in differences in tissue deformation and stabilization and, in selected acquisitions, in glass reflections, signal loss, and motion-related artifacts. The removal of the glass layer in partial detachment areas during preprocessing could eventually remove part of the epidermal signal together with the glass-related artifact, potentially introducing additional variability in the extracted morphological features. These effects can influence both vascular and morphological feature extraction.

The development of a smaller or more adaptable probe interface, capable of maintaining reproducible contact across different anatomical sites, could reduce acquisition-related variability and improve longitudinal comparability. Moreover, future improvements in imaging resolution could help recognizing better those small and disrupted vessels which were not totally defined and limited the automatic segmentation performance.

The OCT morphological processing pipeline also presents methodological limitations. Although the developed procedure enabled automatic identification of the skin surface and DEJ and the extraction of quantitative structural descriptors, segmentation remained challenging in volumes characterized by poor tissue contrast, partial contact, superficial reflections, or poorly defined interfaces. In addition, the current implementation requires relatively long computational times because several processing steps are performed iteratively and have not yet been optimized for large-scale analysis. Future developments could therefore focus both on computational optimization of the existing pipeline and on faster segmentation strategies. In particular, three-dimensional approaches exploiting spatial continuity across the entire volume or learning-based methods could potentially improve robustness and computational efficiency.

The morphological feature set investigated in this thesis represents an initial characterization of the structural information available from OCT and could be further expanded. The implemented descriptors mainly characterize epidermal thickness and skin-surface topography, while additional tissue properties remain unexplored.

Approaches for quantifying regions with poorly defined DEJ, globular dark structures, and stratum corneum thickness could improve the OCT characterization. Their future implementation, together with additional descriptors of tissue texture, layer organization, or optical properties, could provide a more comprehensive characterization of radiation-induced structural changes. This is interesting considering that the present results suggest a complementary contribution of OCT morphology to the vascular and radiomic information.

Moreover, the integration of information about toxicity onset could provide

a more precise characterization of the temporal evolution of tissue response and potentially improve the identification of early toxicity-related changes.

Finally, the clinical heterogeneity of the cohort should also be considered. The dataset included different NMSC and precancerous lesion types, including BCC, SCC, and actinic keratosis, which may present different baseline structural and vascular characteristics and potentially different responses to radiotherapy. The limited sample size did not allow reliable stratified analyses according to lesion type, anatomical location, or timing of toxicity development. Larger datasets could enable these factors to be incorporated explicitly into the analysis, allowing the investigation of lesion-specific treatment responses and potentially improving the interpretation of quantitative OCT/OCTA biomarkers. Ultimately, validation on a larger and independent cohort will be required before the proposed framework can be considered for clinical translation.

4.5 Conclusions

This thesis developed and evaluated a quantitative multimodal OCT/OCTA framework for the non-invasive analysis of skin lesions undergoing superficial radiotherapy. Starting from an existing OCTA-based workflow, the study addressed different aspects of the processing pipeline, including automated vascular segmentation, motion-artifact correction, structural OCT segmentation, and the extraction of morphological descriptors.

The introduction of OCT-derived features extended the analysis beyond vascular characterization, allowing structural changes in the epidermis and skin surface to be investigated throughout treatment and follow-up. Their integration with OCTA-derived and radiomic information provided a more comprehensive representation of tissue response and supported the potential value of combining complementary imaging biomarkers for longitudinal monitoring and radiation-induced toxicity assessment.

In conclusion, this work provides a methodological foundation for a more automated, objective, and multimodal analysis of OCT/OCTA data in patients with non-melanoma skin cancer. Although further validation and methodological refinement are required before clinical translation, the proposed framework represents a step toward quantitative and non-invasive tools capable of supporting the characterization of treatment-related tissue changes. This could potentially improve patient monitoring during radiotherapy.

Bibliography

[1] V. Madan, J. T. Lear, and R.-M. Szeimies. «Non-melanoma skin cancer». In: *The Lancet* 375.9715 (Feb. 2010). DOI: 10.1016/s0140-6736(09)61196-x.

[2] V. Samarasinghe and V. Madan. «Nonmelanoma skin cancer». In: *Journal of Cutaneous and Aesthetic Surgery* 5.1 (Jan. 2012). DOI: 10.4103/0974-2077.94323.

[3] National Cancer Institute. *Common Terminology Criteria for Adverse Events (CTCAE) Version 5.0*. Published: November 27, 2017. U.S. Department of Health and Human Services. 2017. URL: https://dctd.cancer.gov/research/ctep-trials/for-sites/adverse-events/ctcae-v5-5x7.pdf.

[4] A. F. Fercher et al. «Optical coherence tomography - principles and applications». In: *Reports on Progress in Physics* 66.2 (Feb. 2003). DOI: 10.1088/0034-4885/66/2/204.

[5] A. F. Fercher. «Optical coherence tomography – development, principles, applications». In: *Zeitschrift für Medizinische Physik* 20.4 (Dec. 2010). DOI: 10.1016/j.zemedi.2009.11.002.

[6] S. Aumann et al. «Optical Coherence Tomography (OCT): Principle and Technical Realization». In: *High Resolution Imaging in Microscopy and Ophthalmology*. Springer International Publishing, 2019. DOI: 10.1007/978-3-030-16638-0_3.

[7] A. G. Koutsiaris et al. «Optical Coherence Tomography Angiography (OCTA) of the eye: A review on basic principles, advantages, disadvantages and device specifications». In: *Clin. Hemorheol. Microcirc.* (2022), pp. 1–25. DOI: 10.3233/ch-221634.

[8] M. Liu and W. Drexler. «Optical coherence tomography angiography and photoacoustic imaging in dermatology». In: *Photochem. Photobiol. Sci.* 18.5 (2019), pp. 945–962. DOI: 10.1039/c8pp00471d.

[9] G. Rotunno et al. «DERMA-OCTA: A Comprehensive Dataset and Preprocessing Pipeline for Dermatological OCTA Vessel Segmentation». In: *Sci. Data* 12.1 (2025). DOI: 10.1038/s41597-025-05763-6.

[10] D. A. Lintzeri et al. «Epidermal thickness in healthy humans: a systematic review and meta-analysis». In: *Journal of the European Academy of Dermatology and Venereology* (2022). DOI: 10.1111/jdv.18123.

[11] J. Sandby-Møller, T. Poulsen, and H. C. Wulf. «Epidermal Thickness at Different Body Sites: Relationship to Age, Gender, Pigmentation, Blood Content, Skin Type and Smoking Habits». In: *Acta Dermato-Venereologica* 83.6 (2003), pp. 410–413. DOI: 10.1080/00015550310015419.

[12] P. Therkildsen et al. «Epidermal thickness measured by light microscopy: a methodological study». In: *Skin Res. Technol.* 4.4 (1998), pp. 174–179. DOI: 10.1111/j.1600-0846.1998.tb00106.x.

[13] C. Czekalla et al. «Noninvasive Determination of Epidermal and Stratum Corneum Thickness in vivo Using Two-Photon Microscopy and Optical Coherence Tomography: Impact of Body Area, Age, and Gender». In: *Skin Pharmacology and Physiology* 32.3 (2019), pp. 142–150. DOI: 10.1159/000497475.

[14] Ruchir Srivastava et al. «Three-dimensional graph-based skin layer segmentation in optical coherence tomography images for roughness estimation». In: *Biomedical Optics Express* 9.8 (Aug. 2018), pp. 3590–3604. DOI: 10.1364/boe.9.003590.

[15] Francesco Gili. «Quantitative assessment of radiotherapy-induced microvasculature changes in non-melanoma skin cancer using OCTA». Master's thesis. Turin, Italy: Politecnico di Torino, Oct. 2025.

[16] G. Heilemann et al. «Optical Coherence Tomography Angiography (OCTA) Captures Early Micro-Vascular Remodeling in Non-Melanoma Skin Cancer During Superficial Radiotherapy: A Proof-of-Concept Study». In: *Diagnostics* 15.21 (2025), p. 2698. DOI: 10.3390/diagnostics15212698.

[17] C. Blatter et al. «In situ structural and microangiographic assessment of human skin lesions with high-speed OCT». In: *Biomedical Optics Express* 3.10 (2012), pp. 2636–2648. DOI: 10.1364/BOE.3.002636.

[18] J. Weissman, T. Hancewicz, and P. Kaplan. «Optical coherence tomography of skin for measurement of epidermal thickness by shapelet-based image analysis». In: *Opt. Express* 12.23 (2004), p. 5760. DOI: 10.1364/opex.12.005760.

[19] A. Taghavikhalilbad et al. «Semi-automated localization of dermal epidermal junction in optical coherence tomography images of skin». In: *Appl. Opt.* 56.11 (2017), p. 3116. DOI: 10.1364/ao.56.003116.

[20] Annan Li et al. «Epidermal segmentation in high-definition optical coherence tomography». In: *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 2015, pp. 3045–3048. DOI: 10.1109/EMBC.2015.7319034.

[21] C.-H. Lin et al. «Rapid measurement of epidermal thickness in OCT images of skin». In: *Sci. Rep.* 14.1 (2024). DOI: 10.1038/s41598-023-47051-6.

[22] Elena Lucia Errico. «Comparative Analysis of OCTA Reconstruction and Artifact Removal Algorithms for Skin Imaging: Validation on skin-mimicking phantoms and human in-vivo data». Master's thesis. Turin, Italy: Politecnico di Torino, 2026.

[23] Leopold Veselka et al. «A Quantitative Model for Optical Coherence Tomography». In: *Sensors* 21.20 (Oct. 2021), p. 6864. DOI: 10.3390/s21206864.

[24] Y. Zhang et al. «Development of a Novel, Multi-scene Universal Multiple Wavelet-FFT Algorithm (MW-FFTA) for Denoising Motion Artifacts in OCT-Angiography In Vivo Imaging». In: *Optics Express* 30.18 (Aug. 2022), pp. 31735–31752. DOI: 10.1364/oe.465255.

[25] goldbema. *WeightedLeastSquaresFilter: Weighted Least Squares Filter implementation*. Version 1.0. GitHub repository. 2019. URL: https://github.com/goldbema/WeightedLeastSquaresFilter.

[26] P. Lindsø Andersen et al. «Vascular morphology in normal skin studied with dynamic optical coherence tomography». In: *Experimental Dermatology* 27.9 (2018), pp. 966–972. DOI: 10.1111/exd.13680.

[27] B. Heise et al. «Monitoring of Early Skin Reactions by Optical Coherence Tomography and Dermatoscopy in Patients Receiving Radiation Therapy for Head and Neck Cancer». In: *Advances in Radiation Oncology* 10.7 (2025), p. 101793. DOI: 10.1016/j.adro.2025.101793.

[28] W. Jerjes et al. «In Vivo Optical Coherence Tomography for the Detection, Subtyping, and Margin Assessment of Facial Basal Cell Carcinoma: A Comparative Study with Histopathology». In: *Journal of Clinical Medicine* 14.3 (2025), p. 949. DOI: 10.3390/jcm14030949.

[29] R. Srivastava et al. «Three-dimensional graph-based skin layer segmentation in optical coherence tomography images for roughness estimation». In: *Biomedical Optics Express* 9.8 (2018), pp. 3590–3607. DOI: 10.1364/boe.9.003590.

[30] A. P. Yow et al. «Automated in vivo 3D high-definition optical coherence tomography skin analysis system». In: *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2016, pp. 5697–5700. DOI: 10.1109/embc.2016.7591579.

Self-declaration on the use of generative AI

Large Language Models, specifically ChatGPT (OpenAI) and Gemini (Google), were used for suggestions in code implementation and editorial support, such as grammar correction, translations, paraphrasing suggestions, or readability improvements. Regardless of the level of AI involvement, all generated content was critically reviewed, verified, and edited. The author takes full responsibility for the accuracy, originality, and integrity of the final work.