UNIVERSITY OF CALIFORNIA
SANTA CRUZ

COMPUTATIONAL STRATEGIES FOR OPTIMIZING WOUND HEALING: FROM NEURITE ANALYSIS TO DEEP REINFORCEMENT LEARNING

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ELECTRICAL AND COMPUTER ENGINEERING

by

Sam Teymoori

December 2023

The Dissertation of Sam Teymoori
is approved:

Dr. Marcella Gomez, Chair

Dr. Marco Rolandi

Dr. Zouheir Rezki

Peter Biehl
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

List of Tables

Abstract

COMPUTATIONAL STRATEGIES FOR OPTIMIZING WOUND HEALING: FROM NEURITE ANALYSIS TO DEEP REINFORCEMENT LEARNING

by

Sam Teymoori

Abstract The primary objective of this thesis is to expedite the intricate process of wound healing through a multifaceted computational approach. Our journey commences by delving into the relationship between neurites and wound recovery, a relatively unexplored yet pivotal facet of this biological phenomenon.

In the initial phase of our research, we introduce an innovative computational methodology designed to establish correlations between neuronal activity and the progression of wound healing. Employing advanced data analytic techniques, we unravel the pivotal role neurites play in the overall healing process. This newfound understanding significantly enriches our comprehension of wound healing mechanisms.

The second phase of our study harnesses a state-of-the-art deep learning algorithm. This algorithm is engineered to identify and quantify neurites within microscopic images autonomously. This high-throughput approach equips us with the ability to investigate how various pharmaceutical interventions impact neurite growth meticulously. This foundational research lays the cornerstone for the development of optimized strategies aimed at expediting wound healing.

As we progress to the final phase of our investigation, we take a more comprehensive approach. Here, we introduce a Deep Reinforcement Learning (DRL) model, which is tailored to optimize wound healing based on the analysis of wound

images. While the DRL model does not directly incorporate neurite data, the insights gleaned from our earlier studies significantly enhance the model's effectiveness. Trained to recognize various stages of wound development through image analysis, the DRL model adapts in real time, offering a flexible framework for making critical medical decisions.

This thesis, commencing with a focused exploration of neurites and culminating in applying DRL to wound images, signifies a substantial breakthrough in computational medicine. The fusion of deep learning and reinforcement learning techniques presents a potent arsenal for comprehending and enhancing the intricate biological processes underlying wound healing. Ultimately, this research lays the groundwork for more efficient and personalized therapeutic approaches in wound care, promising significant advancements in patient outcomes and healthcare practices.

To my family

whose unwavering belief in me served as a pillar of strength during challenging times. Your encouragement fueled my resolve to persevere and see this journey through to the end.

Acknowledgments

First and foremost, I extend my deepest gratitude to Dr. Marcella Gomez, my academic advisor, whose wisdom, guidance, and unwavering support have made this journey not only achievable but profoundly enriching. Your intellectual generosity transformed every challenge into an opportunity for growth, making my PhD journey a memorable experience.

I would also like to thank Dr. Marco Rolandi for his invaluable mentorship in honing my presentation skills and the art of effective communication. Your relentless encouragement pushed me to continuously refine my work, making a lasting impact on my academic development.

Special thanks are due to Dr. Zouheir Rezki, whose constructive critiques and insightful recommendations substantially elevated the quality of this thesis. Your rigorous academic standards have inspired me to strive for nothing less than excellence.

I owe a debt of gratitude to my teammates who, despite tight deadlines and a fast-paced project environment, collaborated with a spirit of unity and purpose. Your collective effort turned an ambitious project into a successful reality, and I am grateful for each of your contributions.

On a personal note, I wish to acknowledge my friends and family, who provided an unwavering source of emotional support throughout this endeavor. A heartfelt thank you to my parents and my brother, whose enthusiasm for my work often surpassed my own and served as a motivational force.

I must extend a warm acknowledgment to my U.S.-based family and to my friend Tim Morgan. Our many restorative walks have been a source of both comfort

and clarity, helping me navigate the complexities of this research. To any future PhD candidates, I cannot stress enough the therapeutic value of a good friend and long walks.

Part I

First Part

Chapter 1

Introduction

1.1 Background and Motivation

Wound healing is a complex and vital biological process in which the body strives to restore tissue integrity and function following injury. While this process is essential for maintaining health, it often presents challenges due to its multifaceted nature. The traditional approaches to wound care have limitations in terms of effectiveness and personalization, highlighting the need for innovative and data-driven solutions [1].

Recent studies suggest that neurites, projections from nerve cells, may play a role in wound healing. While neurites are well-documented in the context of the nervous system, their function in tissue repair is not extensively researched. This lack of understanding allows computational studies to investigate the relationship between neuronal activity and wound-healing mechanisms [2].

This research is motivated by leveraging advanced computational methodologies to enhance our understanding of wound healing mechanisms and ultimately optimize therapeutic interventions. Within this context, we embark on a multi-phased

investigation to address the complexities of wound healing.

1.2 Identification of Problem

Traditional wound healing approaches are not specific enough to meet patients' needs [3]. Current methods often lead to prolonged healing times and unpredictable results, leaving much room for improvement [4]. Additionally, the potential role of neurites in wound healing is poorly understood, creating a significant knowledge gap. In this context, an alternative control method, such as model predictive control, can be explored. This allows for desired constraints, such as control effort, to be placed.

1.3 Contributions and Objectives

This research is poised to make significant contributions to the field of wound care and computational medicine:

- Enhanced Understanding: By elucidating the role of neurites in wound healing, this research contributes to a deeper understanding of the biological mechanisms at play during the recovery process.
- High-Throughput Analysis: The development of automated neurite identification techniques using deep learning enables efficient analysis, fostering the assessment of pharmaceutical interventions and their impact on neurite growth.
- Dynamic Optimization: Implementing a DRL model for real-time wound healing optimization offers a novel and adaptable approach to wound care, ultimately improving patient outcomes.

The primary objectives of this research are to unravel the complexities of wound healing, bridge the gap between neuroscience and medical imaging, and develop strategies for more efficient and personalized wound care. Through these objectives, we seek to pave the way for a future where advanced computational methods are seamlessly integrated into clinical practice, revolutionizing wound care and improving patients' lives worldwide.

Chapter 2

Previous Work

2.1 The role of neurites in wound healing

Wound healing is a fundamental physiological process, essential for the survival of an organism. It is a complex and dynamic series of events that restores the integrity of damaged tissue following injury. This intricate process is categorized into four distinct but overlapping phases: hemostasis, inflammation, proliferation, and remodeling. Each phase is characterized by specific cellular behaviors and biochemical responses, which must occur in a precise and regulated sequence to ensure successful tissue repair.

The discovery of the Nerve Growth Factor (NGF) marked a significant milestone in neurobiology [5]. Initially identified for its critical role in the growth, maintenance, and survival of nerve cells, NGF has since been recognized as a pivotal player in various biological processes, including wound healing. NGF is a member of the neurotrophin family, a group of proteins that are key regulators of neural development.

NGF's influence extends beyond the nervous system. It is a potent healing

agent, capable of stimulating cell proliferation, differentiation, and survival. In the context of wound healing, NGF serves as a biochemical signal that promotes the repair and regeneration of damaged tissues. It acts on a variety of cell types involved in the healing process, including fibroblasts, keratinocytes, and immune cells. Through its high-affinity interaction with specific receptors, such as TrkA and p75NTR, NGF activates signaling pathways that lead to cellular responses essential for healing.

The role of NGF in wound healing is multifaceted. It enhances angiogenesis, the formation of new blood vessels, which are crucial for supplying nutrients and oxygen to the healing tissue. NGF also modulates the inflammatory response, ensuring it is robust enough to ward off infection but controlled to prevent excessive tissue damage. Furthermore, NGF stimulates the production of extracellular matrix components, facilitating the formation of new tissue and ultimately, scar formation.

The review paper by Zhenxing Liu et al. [6] aims to synthesize the current understanding of NGF's role in wound healing. It examines the molecular and cellular mechanisms by which NGF influences the healing process, and it discusses the therapeutic potential of NGF in treating wounds. By exploring the interplay between NGF and wound healing, the review sheds light on the possibility of developing new treatments that harness the regenerative power of NGF to improve healing outcomes, particularly in chronic wounds where healing is impaired.

As the scientific community continues to unravel the complexities of wound healing, the significance of NGF becomes increasingly apparent. The review not only highlights the therapeutic applications of NGF but also sets the stage for future research that could lead to innovative strategies for managing wounds. The potential of NGF to enhance wound healing represents a promising frontier in regenerative

medicine and underscores the continued relevance of [5] groundbreaking discovery.

Previous investigations into wound healing have consistently underscored the vital role of NGF [7]. NGF interacts with each phase of wound healing in specific and critical ways. Initially, during the hemostasis stage, NGF is involved in the formation of a fibrin clot. This clot acts as a temporary platform, enabling cells to migrate to the wound site and initiate the repair process. As the wound enters the inflammatory phase, NGF plays a role in modulating the body's immune response, attracting essential immune cells that help combat infection and clear away debris [6].

The role of NGF becomes particularly significant during the proliferative phase of healing. Here, NGF stimulates the growth and movement of fibroblasts and keratinocytes, which are central to tissue repair. Fibroblasts are responsible for producing and arranging new matrix materials, while keratinocytes migrate to form a new protective layer over the wound. Additionally, NGF is crucial for angiogenesis in this phase, promoting the development of new blood vessels that supply the necessary nutrients and oxygen to the regenerating tissue [8].

During the final remodeling phase, NGF continues to play a supportive role in the maturation and differentiation of the new tissue. It assists in structuring the new matrix and regulates the process of cell turnover, ensuring the removal of cells that are no longer needed. This phase can be prolonged, but the presence of NGF is essential throughout, as it ensures that the new tissue develops proper strength and functionality [7].

At the molecular level, NGF's function is to bind to a specific receptor that triggers a series of cellular responses. This binding initiates a chain reaction that activates various pathways responsible for cell survival, growth, adaptation, and movement all vital for effective wound healing [6].

NGF also influences another pathway that is implicated in controlling inflammation and immune responses. The precise coordination of NGF's activity ensures that cellular responses are timely and appropriate throughout the healing process [6].

To summarize, NGF's role in wound healing is comprehensive, impacting a broad spectrum of cellular functions and pathways. Its capacity to regulate cell behavior and communication highlights its potential as a therapeutic agent for enhancing wound repair and regeneration. The continued research into the actions of NGF is revealing new ways in which this protein can be leveraged to improve clinical outcomes in wound management.

2.2 The development of neurite identification techniques using deep learning

The development of neurite identification techniques using deep learning has seen significant advancements, particularly with the introduction of the Neurite Annotation Prediction Algorithm (NAPA) [9]. Neurites, the delicate extensions from neuronal cell bodies, are challenging to trace manually due to their thin, dim, and complex nature. Traditional methods of neurite annotation are labor-intensive and subject to limitations in accuracy and scalability, especially when dealing with high-throughput imaging data [10].

Traditional methods of neurite annotation are labor-intensive and limited in their capacity to generalize across various imaging conditions, especially in high-throughput settings like time-lapse or longitudinal imaging. This limitation has necessitated the development of more robust and efficient techniques for neurite

quantification.

NAPA, a deep learning-based neurite annotation prediction algorithm, addresses these challenges by learning broader features crucial for neurite recognition. It utilizes a small number of annotated examples from a dataset to generate annotations for the entire dataset. This approach significantly reduces the time and labor involved in neurite tracing, enabling broad neurite prediction across large and diverse datasets. It also allows for the longitudinal tracking of neurite connections, which is essential for understanding the progression of neurological conditions.

The algorithm was developed using an auto-encoder model, a framework previously employed in applications like image denoising [11]. In this context, the neurite signal is designated as the target feature to learn, while other sources of intensity in the images (such as debris, soma, and imaging artifacts) are considered noise to be removed. The model was trained on datasets with varying degrees of neurite annotation, using binary neurite trace masks paired with original raw images as inputs. This training enabled the model to predict neurite masks with high accuracy, approaching and sometimes exceeding the quality of manual curation.

While deep learning techniques like NAPA represent a significant advancement in neurite identification and analysis, their direct application to wound healing research is limited. The differences in biological processes, imaging data, and specific requirements for algorithm training and interpretation in the context of wound healing present significant challenges that currently hinder the applicability of these techniques in this field.

2.3 Implementing a ordinary differential equation model for real-time wound healing optimization

The exploration of personalized precision treatments in wound healing, particularly through the lens of artificial intelligence (AI), deep learning, and reinforcement learning, represents a significant shift in modern medical practices [12]. This shift is rooted in the understanding that individual patients respond differently to medications due to molecular variations. These variations can occur not only among different individuals but also within the same individual over time. The goal of personalized treatment is to customize drug dosages, types, and the timing of their delivery based on each patient's unique response, utilizing experimental data and statistical analysis for guidance [13].

Historically, wound healing has been understood as a multifaceted process involving several stages: hemostasis, inflammation, proliferation, and maturation [14]. This process is characterized by the nonlinear interaction of different cells (like platelets, neutrophils, macrophages, myofibroblasts, fibroblasts, and keratinocytes) and biomolecules (including blood coagulation factors and various cytokines). The complexity of this process has traditionally posed a challenge in determining the optimal timing and dosage for drug administration, especially considering the nonlinear dynamics of drug absorption and the specific biological targets of each drug [15].

In response to these challenges, researchers have been developing innovative methods to improve the precision and effectiveness of wound healing treatments [16,17]. One such approach is the use of online self-adaptive controller methods that employ deep learning and reinforcement learning, tailored to the patient's real-time response to treatments. However, these methods have not yet been extensively explored by

researchers.

The development of these advanced treatment methods has been influenced by the need to account for individual variability in pharmacokinetics and pharmacodynamics. Factors such as gender, age, weight, and overall health conditions can significantly impact how a patient responds to a particular treatment [18]. Reinforcement learning-based approaches are particularly advantageous in this context, as they allow for the development of adaptable control algorithms that can optimize drug dosing in real-time, even in the presence of patient variability [19].

These learning-based control strategies have found applications in various medical settings, enhancing the precision of drug dosing and optimizing treatment regimens. However, a significant limitation of these approaches is their lack of consideration for safety exploration in reinforcement learning. The inherent complexity of deep neural networks and the non-convex nature of objective functions make it challenging to achieve a globally optimal control policy [20]. The challenge in this area of research lies in finding the optimal mapping between nonlinear systems and linear models.

In summary, the field of wound healing stands to gain significantly from the integration of AI, deep learning, and reinforcement learning techniques. These methods are propelled by the necessity for personalized precision treatments that cater to the unique variability of individual patients. Despite facing challenges in ensuring the safety and optimality of treatment protocols, the research and development in this domain have the potential to make a substantial impact on the advancement of personalized medicine.

Chapter 3

Quantifying innervation facilitated by deep learning in wound healing

This chapter is from [21] as it appears in Scientific Reports. The dissertation author was a co author of this paper.

Abstract

The peripheral nerves (PNs) innervate the dermis and epidermis, and are suggested to play an important role in wound healing. Several methods to quantify skin innervation during wound healing have been reported. Those usually require multiple observers, are complex and labor-intensive, and the noise/background associated with the immunohistochemistry (IHC) images could cause quantification errors/user bias. In this study, we employed the state-of-the-art deep neural network, Denoising Convolutional Neural Network (DnCNN), to perform pre-processing and effectively reduce the noise in the IHC images. Additionally, we utilized an automated image analysis tool, assisted by Matlab, to accurately determine the extent of skin innervation during various stages of wound healing. The 8 mm wound is generated using a circular biopsy punch in the wild-type mouse. Skin samples were collected

on days 3, 7, 10 and 15, and sections from paraffin-embedded tissues were stained against pan-neuronal marker- protein-gene-product 9.5 (PGP 9.5) antibody. On day 3 and day 7, negligible nerve fibers were present throughout the wound with few only on the lateral boundaries of the wound. On day 10, a slight increase in nerve fiber density appeared, which significantly increased on day 15. Importantly, we found a positive correlation ($R = 0.96$) between nerve fiber density and re-epithelization, suggesting an association between re-innervation and re-epithelization. These results established a quantitative time course of re-innervation in wound healing, and the automated image analysis method offers a novel and useful tool to facilitate the quantification of innervation in the skin and other tissues.

3.1 Introduction

Wound regeneration is a complex process that is regulated by orchestrated mechanisms, influenced by chemical, cellular, and molecular factors [22, 23]. The healing process begins at the time of injury and eventual maturation could continue for months or even years until the wound completely heals and is structurally and functionally similar to uninjured skin [24]. The four overlapping phases of wound healing are homeostatic, inflammatory, proliferative, and remodeling. The homeostatic phase lasts a few hours, producing a fibrin plug followed by an inflammatory phase, which can last between hours and days, during which aggregated platelets and cells release pro-inflammatory mediators [25]. The early inflammatory phase is succeeded by the proliferative phase lasting a few weeks during which macrophages and fibroblast cells invade the wound bed forming granulation tissue and active migration of the wound epithelial cells occurs [26]. The last phase of wound healing is the

remodeling phase, which is characterized by proliferative cell apoptosis, adjustment of extracellular matrix (ECM), and replacement of type 3 with type 1 collagen in the dermis. It can last for weeks to years [24].

During healing phases, a strong interaction between the nervous system and skin involving a variety of neuromodulators, cytokines, hormones, and other effector molecules has been reported [27, 28]. The nervous system can be influenced both at the local and central levels by the stimuli at the skin and vice versa. The brain can alter skin function during the pathophysiological state and skin can modulate the nervous system by releasing a variety of neuropeptides [27, 29]. The cutaneous nerves therefore positively affect all the stages of wound healing [30]. Numerous neuropeptides e.g. substance P (SP) released from cutaneous nerves have been reported to activate vital mechanisms during the inflammatory phase [31]. Similarly, neuropeptides released by cutaneous nerves influence the proliferation phase. They can promote the proliferation of fibroblasts, keratinocytes, and endothelial cells by stimulating DNA synthesis, can stimulate angiogenesis, support granulation tissue remodeling, and many more [26, 32–35]. The effect of innervation on the remodeling stage has also been studied in the past. It has been demonstrated that a significantly higher number of nerve fibers correlate with normotrophic scars in comparison to hypertrophic scars during the remodeling phase [36,37]. Hence, the literature strongly suggests a regulatory role for skin nerves in wound healing and any impairment in skin innervation is one of the leading causes of occurrence of chronic wounds e.g., diabetic neuropathy can lead to foot ulcers and plegias to sacral and trochanteric pressure sores [38, 39].

Previously, numerous studies have been conducted on quantifying skin innervation [40–48]. However, most such studies are not fully automated, have

manual counts of IHC-stained structures that are prone to user errors and variation, require multiple observers, and are complex and labor-intensive. One such example is dendrite analysis involving manually tracing neurons using the simple neurite tracer plug-in of ImageJ software [22, 42]. While this semi-manual approach has been proven effective, it involves identifying the beginning and end points of dendrites and digitally drawing individual branch segments throughout the entire neuron, making it a labor-intensive and time-consuming process. Another example is semi-automated Sholl analysis for quantifying changes in the growth and differentiation of neurons and glia [43]. The method offers several advantages over conventional manual quantification, including faster analysis time and increased statistical sensitivity. However, the method has some limitations, such as reliance on manual input from the user, which introduces a risk of user error and variation impacting the accuracy and reliability of the results. Additionally, the semi-automated Sholl method is complex and time-consuming to set up initially, which could act as a barrier for researchers who do not have the technical expertise or resources to implement the method effectively. In an effort to quickly, objectively, and reproducibly quantify cutaneous innervation, we developed a fully automated Matlab-assisted image analysis tool aided by the deep neural network, DnCNN, for pre-processing (de-noising) of the IHC-images. This network can detect and remove high-frequency image artifacts and increase image resolution. Image noise is minimized, resulting in higher-quality images that can be more accurately analyzed. The DnCNN is particularly developed for image processing [49], and has shown effectiveness in a wide range of applications, including medical imaging [50].

Utilizing an automated Matlab-assisted tool aided with DnCNN we quantified skin innervation during wound healing stages at days 3, 7, 10 and 15 (Figure

3.1). The data show a positive correlation between the increase in nerve fiber density and re-epithelization.

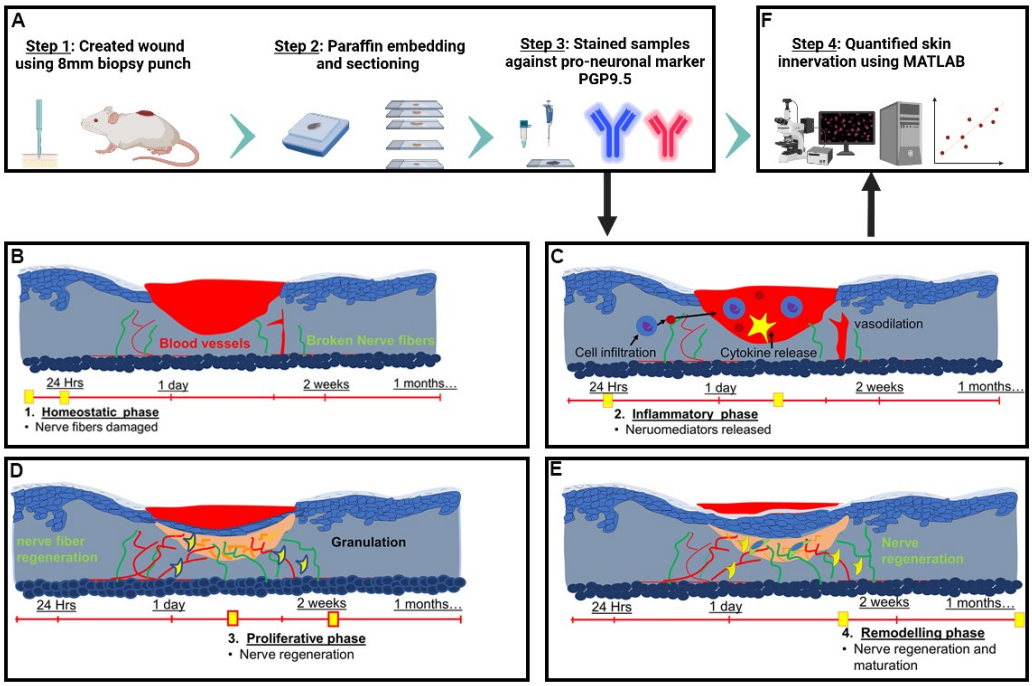

Figure 3.1: The experimental design and schematic depicting the methodology used to quantify skin innervation. (A) A biopsy punch of 8 mm in diameter is used to create the wound, and skin samples are collected and fixed on days 3, 7, 10 and 15. After fixation, the wounded tissue is paraffin-embedded and sectioned ($5\mu m$ thickness) for immunofluorescence analysis against PGP9.5 protein, a neuron-specific marker. (B–E) Illustration portraying different stages of wound healing. (B) The homeostatic phase lasts a few hours, during which nerve fibers in the wound bed are damaged, followed by the (C) inflammatory phase, which can last between hours and days. (D) The proliferative phase lasts a few weeks, during which re-innervation might be initiated, and (E) during the remodeling phase wound matures and can last between weeks to years. In our study, we chose to quantify skin innervation at days 3, 7, 10 and 15 as an attempt to cover all phases of wound healing. (F) The immunohistochemistry (IHC) samples are analyzed using automated Matlab-assisted tools aided by DnCNN-based image denoising. The images were created with BioRender.com.

3.2 Materials and Methods

3.2.1 Animals

All methods were approved and performed in accordance with the relevant guidelines and regulations of the Institutional Animal Care and Use Committee at UC Davis, were reviewed and approved by the institution's IACUC and performed at the UC Davis Teaching and Research Animal Care Services (TRACS) facility. Male C57BL/6 (older than 28 weeks, 30–35 g) mice were obtained from Jackson Laboratory. The animals were acclimated for one week after transfer from the vendor to the UC Davis vivarium. All male mice used in this experiment were kept in the containment unit of the animal facility, housed in cages with free access to food and water. Mice would be excluded from the experiment when they were at the anagen phase of hair cycle due to variation of healing rate, or when they lost $> 20\%$ weight during the experimental period. The animals were randomly assigned to the treatment groups based on their weight prior to the surgery. All experiments were performed in a biosafety cabinet and were done in triplicate (n=3). The study is reported in accordance with ARRIVE guidelines.

3.2.2 Wounding

The animals were sedated with 3–5% isoflurane two days prior to wounding surgery and the dorsal surface was shaved. To remove the remaining hair depilatory cream was applied and removed within 5–10 s of application. For surgery, all mice were anesthetized with isoflurane. Buprenorphine (0.05 mg/kg) was administered subcutaneously prior to the wounding procedure. Iodine and ethanol wipes were used to sterilely prep experimental mice. An 8 mm sterile skin biopsy punch instrument was

used to create full-thickness wounds on the dorsal skin as described previously [28,51]. In this study, we used a splinted wound model to avoid wound contraction [52,53]. At the end of the experiment, mice were euthanized by cervical dislocation under 5% isoflurane, and skin samples were collected on days 3, 7, 10, and 15 post-wounding.

3.2.3 Immunohistochemical analysis of wound tissues

After fixation, the wound tissue was paraffin-embedded and sectioned (5 μm thickness) for immunohistochemical analysis as described previously [54,55]. Primary antibodies against PGP9.5 (Invitrogen, Catalog # PA5-29012), and β-III tubulin (Invitrogen, Catalog # MA5-16308) were used. Donkey anti-Rabbit IgG (H+L) highly Cross-Adsorbed Secondary Antibody, Alexa Fluor™ 594 (Invitrogen, Catalog # A21207), and donkey anti-Mouse IgG (H+L) Highly Cross-Adsorbed Secondary Antibody, Alexa Fluor™ 488 (Invitrogen, Catalog # A21202), respectively were used. VECTASHIELD® Antifade Mounting Media with DAPI (Vector Laboratories, Catalog # H-1200-10) was used to stain nuclei. Slides were imaged with an Olympus FV3000 Confocal Laser Scanning Microscope (Shinjuku City, Tokyo, Japan) as described previously [56–58] and analyzed using Matlab 2021 Image processing.

3.2.4 Quantification

In order to accurately quantify wound healing, several steps are taken to ensure precision and accuracy in counting PGP9.5 positively stained pixels. One of the most important steps is preprocessing the images to remove any unwanted noise or artifacts that could potentially interfere with the analysis. Many factors can contribute to image noise, including non-specific staining, autofluorescence, equipment malfunctions, motion blur, and environmental factors. To eliminate noise from the

images, a deep neural network known as the DnCNN network is employed. It is a convolutional neural network (CNN) specifically designed for image denoising [49]. It works by learning to map between noisy images and clean images and then use this mapping to remove noise from new images. The DnCNN network is trained on a large dataset of noisy and clean images and has shown effectiveness in a wide range of applications, including medical imaging [50].

The DnCNN employed in this study is a trained model derived from the deep network training in MATLAB. Utilizing the Deep Learning Toolbox's pretrained DnCNN model, we integrated it within the MATLAB environment and invoked it through the MATLAB Deep Learning Toolbox. Our process involved denoising images by employing the DnCNN model from the Deep Learning Toolbox, ultimately producing denoised outputs. DnCNN effectively addresses this by enhancing capacity and flexibility for exploiting image attributes. Moreover, it can accelerate training and enhance denoising outcomes. Additionally, DnCNN, grounded in trained data, can discern pixels to filter and retain, thereby eliminating the need for manual adjustments across diverse images.

Once the images have been denoised, the next step is to count the PGP9.5 positive pixels in the images. Positive pixels representing innervation were identified by their immunoreactivity to PGP9.5 antibody (red color). However, identifying those pixels can be challenging because of the variation in the intensity of fluorescence due to different wound depths, healing stages, expression levels and sizes of expression areas. Additionally, the background intensity can pollute the pixel value counted as nerve fibers and nerve terminals. To overcome these challenges, we calculate the background value. This is achieved by identifying the low spectrum values of pixels and removing outliers. The background value serves as a reference point for

subsequent analyses. By subtracting the background value from all pixels in the image, we effectively normalize the pixel values, making them more comparable across different images. This normalization step is essential for accurate comparisons and ensures that variations in background intensities do not influence our results. Next, we employ a robust method for identifying and omitting outliers from the image. We utilize a percentile-based approach where we focus on the 5 percent of data points that have the longest distance from the median. This 5 percent threshold has been carefully adjusted based on the specific characteristics of the image under analysis. It allows us to address the unique properties and variations in pixel intensity that may exist in different types of images. If a pixel value falls within this 5 percent range, it is considered an outlier and is excluded from further analysis. With the fluorescence threshold in place, we proceed to count the number of PGP9.5-positive pixels in the image. Each pixel's (red fluorescence value) is compared to the threshold, and any pixel with a value above the threshold is considered a positive neurite containing pixel count. This process is repeated for all the pixels in the image, determined together as the total number of neurites containing pixels. Next, to identify nerve fiber density in a target region of interest we divided the total number of positive pixels by the corresponding area. In addition to measuring nerve fiber density in the whole wound, we also analyzed the density of nerve fibers (positive pixels) separately for both the epidermis and dermis to get a spatial understanding.

3.2.5 Reepithelization

After fixation, the wound tissue was paraffin-embedded, sectioned to 5 μm and H&E stained for determination of wound re-epithelialization as described previously [51, 55]. Briefly, the BioRevo BZ-9000 inverted microscope (Keyence,

Osaka, Japan) was used to image all the histological sections. Measurements were done by an investigator blinded to experiment with the BZ-II viewer and analyzer (Keyence, Japan). The absence of underlying adipose tissue and hair follicles defines wound edges and wound healing is determined by the re-epithelialization of the epidermis layer [53, 59]. The outgrowth of the newly formed epidermis was tracked manually from the wound edges and the percentage of the combined length of the re-epithelialization to the total length of the wounds was calculated.

3.2.6 Statistics

Statistical analysis was performed using an unpaired, two-tailed student t-test as described previously [60, 61]. Data are expressed as mean ±SD. A P-value less than 0.05 was considered statistically significant.

3.3 Results

3.3.1 PGP9.5 as a specific neuronal marker

Anti-PGP9.5 antibody labels UCHL1/PGP 9.5 protein in the tissue sections. The protein is highly conserved and localized in neurons and is considered as a pan-neuronal marker that stains both sensory as well as autonomic nerves [62, 63]. To deduct background signal, we carried out experiments with negative control for every sample, omitting the primary antibody (Figure 3.2). The background signals from the negative control are determined and used for generating a cutoff window for quantification of the true PGP9.5+ signals during MATLAB-based quantification.

PGP9.5 staining is used widely and considered the gold standard for labeling skin innervation [47, 48, 62, 64], so throughout the study, we quantified PGP9.5

immunofluorescence staining for the detection of nerve fibers.

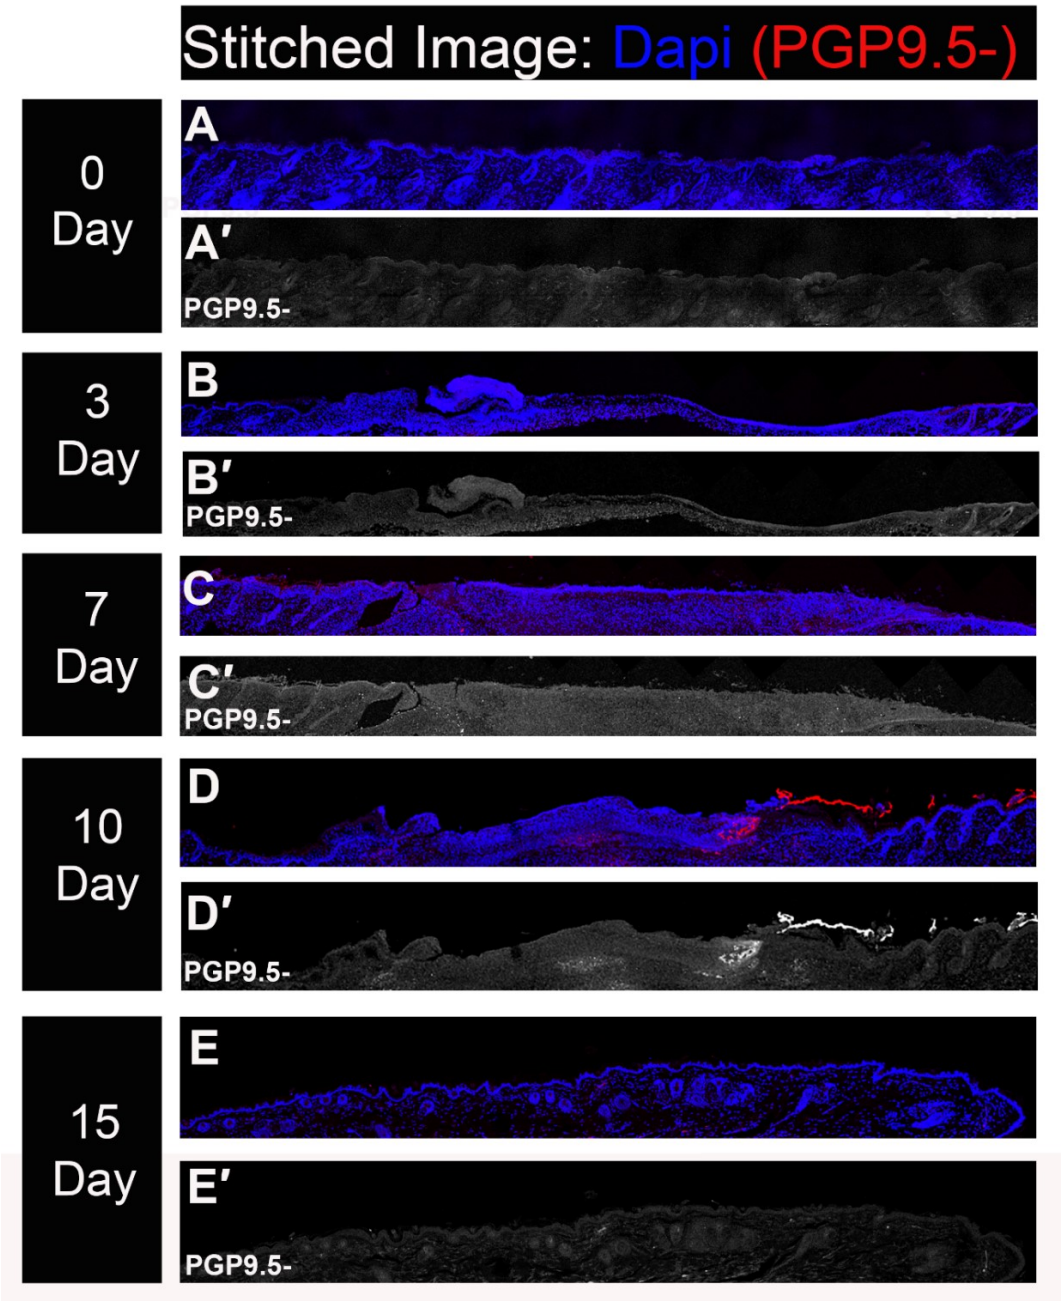

Figure 3.2: Stitched Immunohistochemistry images of 5 μm vertical sections of punch biopsies as a negative control for PGP9.5. PGP9.5 is a pan-neuronal marker and DAPI stains the nuclei (In blue). (A) Uninjured skin. Skin samples were collected on (B) day 3 (C) day 7 (D) day 10 (E) day 15. The red color is autofluorescence from a dead skin flap.

3.3.2 Automated image processing and quantification of skin innervation

The quantification method used in this study employed an automated Matlab-assisted image analysis tool with a deep neural network to pre-process images and determine the range of skin innervation during different stages of wound healing. The use of the deep neural network ensured that noisy pixels did not affect the calculation of total neuronal coverage (Figure 3.3). The depicted layers present how DnCNN operates, transforming a noisy wound image into a clearer and diagnostically valuable representation. Each layer embodies a distinct phase of noise reduction and feature enhancement, showcasing the model's capacity to discern between essential wound structures and undesirable noise. The first layer of the DnCNN model is responsible for capturing the low-level features of the image. During this stage, the initial noise present in the wound image is identified and basic filters are applied to smooth out the noise while preserving the main structures of the wound. As the image progresses through the second layer, which is part of the intermediate layers of the DnCNN model, more complex features and patterns are recognized, contributing to enhanced noise reduction and structure preservation. Each subsequent layer refines the denoising process, gradually enhancing the clarity of the image. The last layer of the DnCNN model is designed to further refine the denoised image and restore it closer to its original state. By this stage, most of the noise has been removed, and wound structures are significantly clearer. The final layer ensures that the denoised image maintains its clinical relevance by minimizing any artifacts introduced during the denoising process, thereby preserving the fidelity of the wound's clinical representation [49, 65, 66].

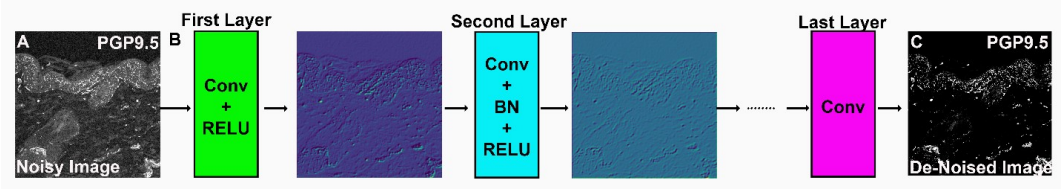

Figure 3.3: DnCNN network architecture for image denoising. (A) Noisy image as DnCNN input. (B) The DnCNN network architecture consists of multiple convolutional layers. Each convolutional layer includes batch normalization (BN), convolution (Conv), and rectified linear unit (ReLU) layers. The first layer takes the noisy image as an input, and the subsequent layers process the image to remove noise. (C) Output image after de-noising.

Throughout the layers, the DnCNN model employs a combination of convolutional filters and nonlinear activation functions to transform the image and remove noise while preserving important features. The progressive nature of the layers enables the model to iteratively refine the denoising process, resulting in a cleaner and clearer wound image that can aid in accurate diagnosis and assessment [58,67,68]. Keep in mind that the specific appearance and effectiveness of the images after passing through each layer depend on the model architecture, training data, and noise characteristics. It is recommended to visually analyze the images to understand the improvements brought about by each layer.

In addition to visual analysis, assessing the Signal-to-Noise Ratio (SNR) improvement at each stage is vital for a quantitative evaluation of the model's performance. By measuring the SNR before and after the image passes through the DnCNN layers, one can quantitatively determine the extent of noise reduction and clarity enhancement. This metric not only validates the perceptual observations but also provides a standardized way to compare the effectiveness of different model

configurations or training approaches in noise reduction.

Once the image has been processed by the DnCnn, our next step is to accurately identify the neurites using a statistical approach. If a threshold is too high, some neurites might not be detected, resulting in an underestimation of the neuronal coverage. Conversely, if the threshold is too low, non-neuronal elements in the image may be mistakenly identified as neurites, leading to an overestimation of neuronal coverage. To address this issue, we employed a statistical approach to determine the appropriate threshold value. First, we saved the R values of every pixel in the image to a list. We then calculated several statistical measures, including the minimum, maximum, mean, median, and interquartile range. Outliers were identified based on this statistical analysis and excluded from further consideration. To determine whether a pixel contained a neurite or not, we used the distance of the R-value from the third interquartile (Q3). Pixels with a closer distance to the maximum value in comparison to Q3 were selected as neurites. This method ensured that the threshold value was based on a statistically robust approach, which increased the accuracy of the quantification method.

Overall, the automated approach demonstrated in this paper required minimal manual intervention, making it suitable for large-scale studies and a reliable method for quantifying skin innervation during wound healing. Additionally, to demonstrate the significance of DnCNN and its impact on noise reduction, we compared two denoising techniques: the conventional threshold denoiser and the advanced DnCNN model (Figure 3.4). Using only conventional threshold denoiser, the nerve fiber density for uninjured skin (unwounded) at outer edge 1 of the wound is found to be 0.04 $\pm$ 0.003 pixels/mm^2 with a density of 0.08$\pm$0.03 pixels/mm^2 found in the epidermis and 0.03$\pm$0.007 pixels/mm^2 in the dermis. Comparatively

incorporating the DnCNN model the nerve fiber density for the same set of images is found to be 0.22 ± 0.03 pixels/mm2 with a density of 0.4±0.1pixels/mm^2 in the epidermis and 0.08±0.02 pixels/mm^2 in the dermis. The conventional threshold denoiser methods employ preset threshold values to detect and filter pixel noise and these approaches struggle to differentiate true noise pixels from low-value pixels, leading to potential misclassification. Further, setting threshold values for each image proves cumbersome. To minimize noise, we adjusted the threshold iteratively until background noise became imperceptible. As a result, the threshold denoiser tends to inaccurately remove neurite positive pixels due to its indiscriminate approach, resulting in underestimating the neurite pixel density. Conversely, the DnCNN required no such adjustment as the DnCNN, grounded in trained data, learns to discern pixels to filter and retain them (Figure 3.4). This highlights the advantage of the DnCNN in accurately maintaining crucial diagnostic information in wound images.

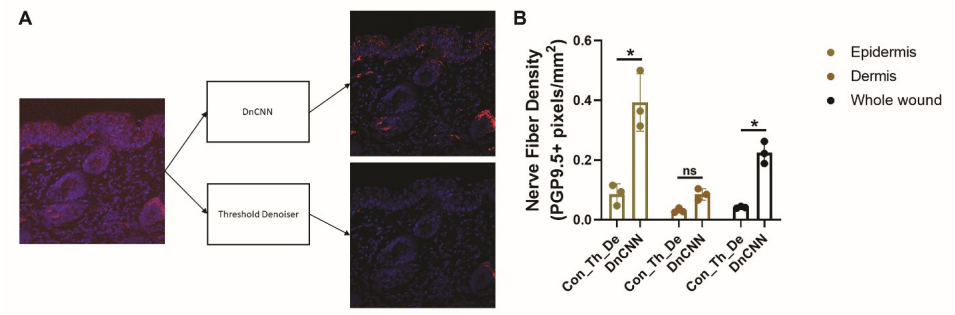

Figure 3.4: Comparison between conventional threshold denoiser and the advanced DnCNN model. (A) PGP9.5 staining (in red) at outer_edge 1 of uninjured skin and subsequent denoising of the stained image using the DnCNN model compared to threshold denoiser, (B) Quantification of innervation in the epidermis, dermis, and whole wound of uninjured skin outer_edge 1 using conventional threshold denoiser (Con_Th_De) and advanced DnCNN model. All quantification data are represented as mean ± SD, $n = 3$ wounds from three mice in each group, $^{*}P < 0.05$, ns = non-significant.

Furthermore, we tested our automated method to analyze neurite innervation for 30 µm-thick cross-sections. The nerve fiber density at the dermis, epidermis and whole wound together showed non-significant change on comparing 5 µm-thick cross-sections to 30 µm-thick cross-section (Figure 3.5). Therefore, results clearly demonstrate that the algorithm used in this study is efficient enough to detect neurites in a 30 µm-thick cross-section of uninjured skin as well, where whole nerve fibers are easier to visualize, but the background can blur the quantification.

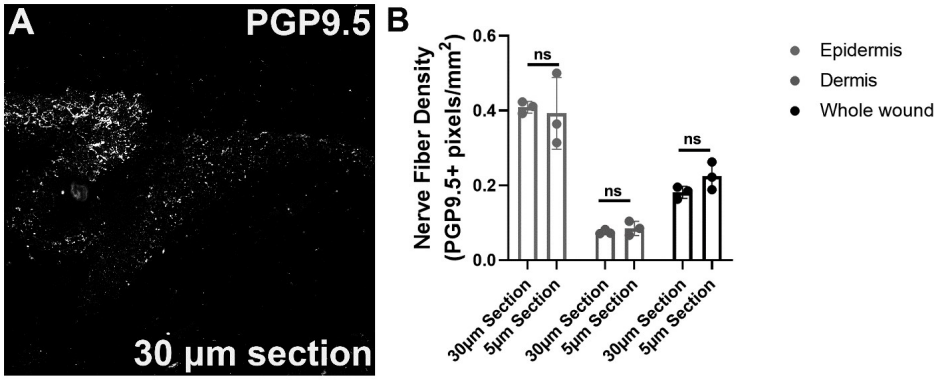

Figure 3.5: Comparison between 5μm thick and 30μm thick section. (A) PGP9.5 staining at outer_edge 1 of uninjured skin for 30μm thick section. (B) Quantification of innervation for 5μm and 30μm thick section in epidermis, dermis and whole wound at outer_edge 1 of uninjured skin. All quantification data are represented as mean $\pm$ SD, $n = 3$ wounds from three mice in each group, ns = non-significant.

3.3.3 Wounding reduced nerve fiber density

Immunoreactivity was detected in intraepidermal and dermal nerve fibers and cells. The positively stained nerve fibers are quantified for the whole wound bed (Figure 3.6; Table 3.1) and separately for the lateral wound boundaries and the wound center (Figure 3.8; Table 3.1). The nerve fiber density individually for the epidermis and dermis is also quantified, respectively (Figure 3.7; Table 3.1). The nerve fiber density for uninjured skin (unwounded) is found to be 0.29 $\pm$ 0.07 pixels/mm^2, for which a density of 0.64$\pm$0.21 pixels/mm^2 is found in the epidermis and 0.27$\pm$0.07 pixels/mm^2 in the dermis. Compared to uninjured skin, as expected after creating a wound, there is a considerable reduction in nerve fiber density throughout the wound bed (WB) (Figure 3.6; Table 3.1).

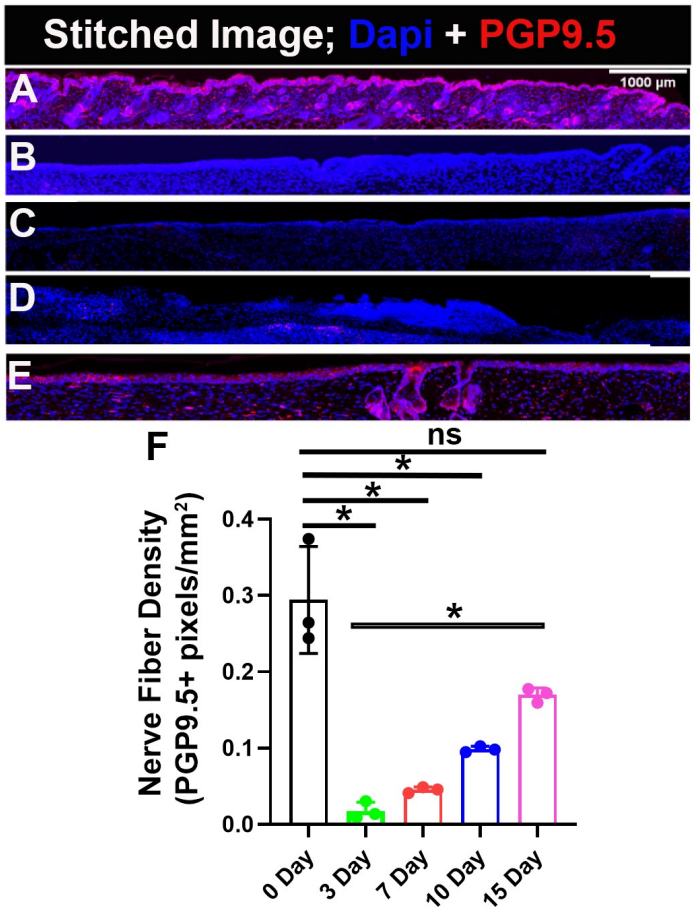

Figure 3.6: Gradual increase in Re-innervation in the wound bed. PGP9.5 is a pan-neuronal marker and DAPI stains nuclei. (A) Uninjured skin. Skin samples were collected on (B) day 3, (C) day 7, (D) day 10, and (E) day 15 of wound healing. (F) Quantification of skin innervation for the whole wound bed represented as mean±SD, n=3 wounds from three mice in each group, *P$<$ 0.05, ns is non-significant. The wound bed is recognized by the absence of hair follicles. Scale bar=1000 μm.

Table 3.1: Nerve fiber density (pixels/mm^2) at wound bed, wound outer edge 1, wound center, and wound outer edge 2 on day 0, 3, 7, 10, and 15 of healing.

Epidermis					Dermis					Whole wound				
0 Day	3 Day	7 Day	10 Day	15 Day	0 Day	3 Day	7 Day	10 Day	15 Day	0 Day	3 Day	7 Day	10 Day	15 Day
Wound bed														
0.64 ± 0.21	0.0 ± 0.0	0.0 ± 0.0	0.08 ± 0.004	0.44 ± 0.04	0.27 ± 0.07	0.02 ± 0.01	0.04 ± 0.003	0.10 ± 0.004	0.16 ± 0.007	0.29 ± 0.07	0.02 ± 0.01	0.045 ± 0.004	0.098 ± 0.004	0.169 ± 0.009
Wound outer edge 1														
0.4 ± 0.1	0.01 ± 0.01	0.04 ± 0.01	0.1 ± 0.04	0.2 ± 0.02	0.08 ± 0.02	0.004 ± 0.001	0.01 ± 0.001	0.01 ± 0.002	0.05 ± 0.003	0.22 ± 0.04	0.01 ± 0.003	0.023 ± 0.005	0.06 ± 0.024	0.12 ± 0.008
Wound center														
0.2 ± 0.02	0 ± 0	0 ± 0	0 ± 0	0.09 ± 0.02	0.13 ± 0.02	0.003 ± 0.001	0.009 ± 0.003	0.016 ± 0.001	0.066 ± 0.01	0.132 ± 0.021	0.003 ± 0.001	0.008 ± 0.003	0.015 ± 0.001	0.07 ± 0.01
Wound outer edge 2														
0.28 ± 0.07	0.03 ± 0.004	0.06 ± 0.006	0.15 ± 0.011	0.243 ± 0.018	0.13 ± 0.01	0.006 ± 0.002	0.009 ± 0.004	0.04 ± 0.01	0.10 ± 0.003	0.19 ± 0.025	0.01 ± 0.003	0.03 ± 0.005	0.092 ± 0.0085	0.16 ± 0.009

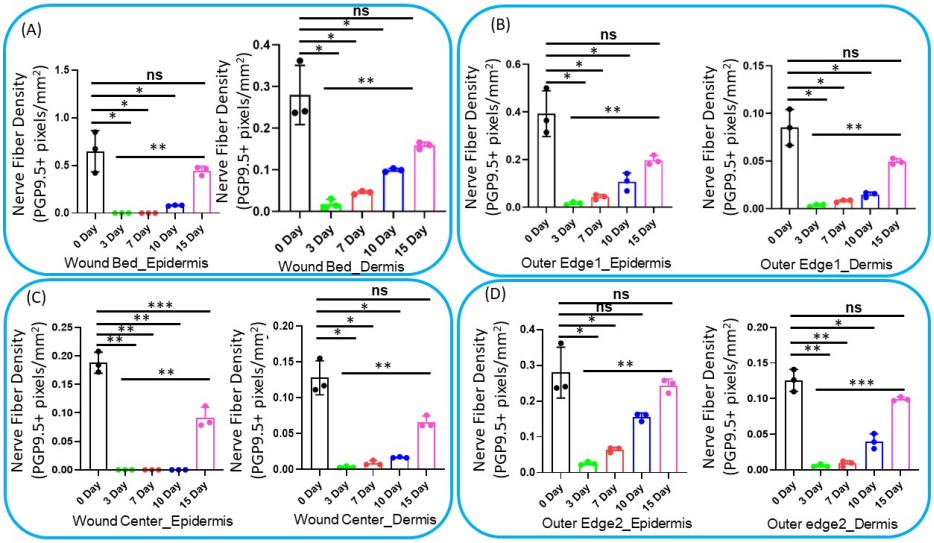

Figure 3.7: Intraepidermal and dermis innervation. Quantification of innervation in epidermis and dermis for (A) Wound bed, (B) Wound outer edge 1, (C) Wound center, (D) Wound outer edge 2. All quantification data are represented as mean ± SD, $n = 3$ wounds from three mice in each group, $^{*}P < .05$, $^{**}P < .001$, $^{***}P < .0001$, ns is non-significant.

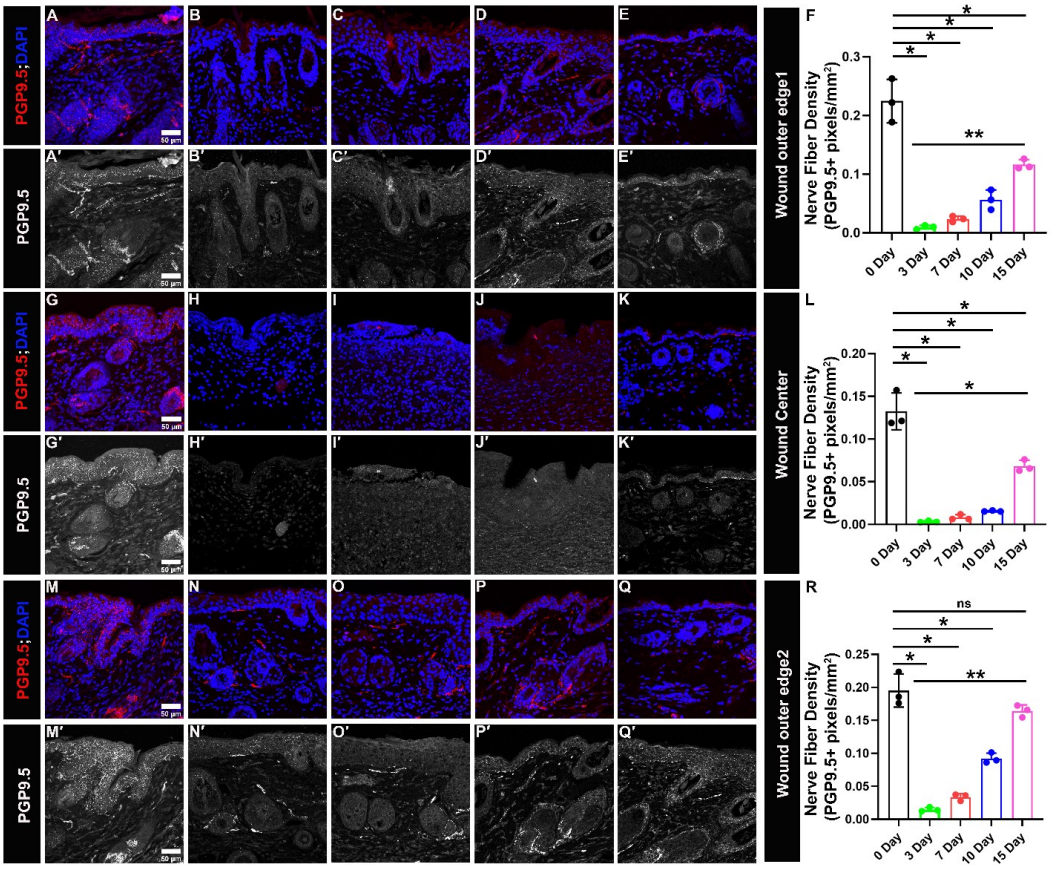

Figure 3.8: Gradual increase in Re-innervation at lateral wound edges and wound center. PGP9.5 is a pan-neuronal marker (in red) and DAPI stains nuclei of cells (in blue). (A'–E', G'–K', M'–Q') Split images (in grey) show PGP9.5 immunoreactivity. (A–E and A'–E') Immunoreactivity to PGP9.5 at wound outer edge 1 for skin samples. (A,A') Uninjured, (B,B') day 3, (C,C') day 7, (D,D') day 10, (E,E') day 15. (F) Quantification of skin innervation at wound outer edge 1. (G–K and G'–K') Immunoreactivity to PGP9.5 at wound center for skin samples. (G,G') Uninjured, (H, H') day 3, (I,I') day 7, (J,J') day 10, (K,K') day 15. (L) Quantification of skin innervation at wound center. (M–Q and M'–Q') Immunoreactivity to PGP9.5 at wound outer edge 2 for skin samples. (M,M') Uninjured, (N,N') day 3, (O,O') day 7, (P,P') day 10, (Q,Q') day 15. (R) Quantification of skin innervation at wound outer edge 2. All quantification data are represented as mean±SD, n=3 wounds from three mice in each group, *P$<$ 0.05, **P$<$ 0.001, ns is non-significant. Scale bar=50 μm.

3.3.4 Innervation gradually appears and reaches a level close to intact skin—15 days after wounding

On day 3 of wound healing, nerve fiber density in the wound bed is found to be 0.02±0.01 pixels/mm^2, which is significantly less than the unwounded skin (0.29 ± 0.07 pixels/mm^2). The nerve fiber density reported (0.02±0.01 pixels/mm^2) is found only in the dermis, yet the epidermis has not regenerated in the wound bed, so the intraepidermal nerve fiber density value is 0. On day 7, the nerve fiber density increases slightly to 0.045±0.004 pixels/mm^2. The intraepidermal nerve fiber density is still 0. The wound sections on day 10 show some traces of intraepidermal nerve fibers; the nerve fiber density quantified is 0.080±0.005 pixels/mm^2. The nerve fiber density in the dermis, 0.10±0.004 pixels/mm^2, also steadily shows an increasing trend. The nerve fiber density for the whole wound bed is 0.098 ± 0.004 pixels/mm^2 on day 10, which is still significantly less compared to unwounded skin (0.29 ± 0.07 pixels/mm^2). Interestingly, on day 15 of healing, there was a huge increase in intraepidermal nerve fiber density, 0.44±0.04 pixels/mm^2. The nerve fiber density in the dermis also substantially increases to 0.16±0.007 pixels/mm^2. On day 15, the changes in nerve fiber density for the whole wound bed compared to unwounded skin (0.17 ± 0.009 pixels/mm^2 vs 0.29 ± 0.07 pixels/mm^2) are non-significant (Figures 3.6 and 3.7) and interestingly significant when compared to day 3 of healing (0.17 ± 0.009 pixels/mm^2 vs 0.02±0.01 pixels/mm^2). The data narrates that there is a gradual increase in nerve fiber density through the time series of wound healing, and there is a significant re-innervation and values start reaching close to normal from day 15 of healing onwards.

3.3.5 The wound center and wound edge show a similar trend in re-innervation

We also quantified nerve fibers separately for wound edges and wound center on days 3, 7, 10, and 15 of wound healing and compared it with the innervation of unwounded skin (Figures 3.8 and 3.7; Table 3.1). For unwounded skin, the density values obtained at outer edge 1 are: intraepidermal (IE)=0.4±0.1 pixels/mm^2; dermis (D)=0.08±0.02 pixels/mm^2; whole unwounded skin=0.22 ± 0.03 pixels/mm^2. It is evident that on days 3, 7 and 10, there is a significant decrease in the IE and D nerve fiber density compared to uninjured skin (Table 3.1). As expected, the nerve fiber density steeply increased on day 15 of healing: IE=0.2±0.02 pixels/mm^2; D=0.05±0.003 pixels/mm^2; whole wound outer edge 1=0.12 ± 0.008 pixels/mm^2. The values for IE and D show non-significant change compared to uninjured skin except for whole wound outer edge 1 where nerve fiber density is still significantly less compared to uninjured skin (0.12 ± 0.008 pixels/mm^2 vs 0.22 ± 0.03 pixels/mm^2). Also, the change found is significant compared to 3 days for IE, D and whole wound outer edge 1 (Figures 3.8 and 3.7; Table 3.1). Therefore, data signify that at the outer edge 1 of the wound, a considerable level of re-innervation takes place up to day 15 of healing. However, total nerve fiber density cannot reach the normal level i.e. the values are significantly less than the uninjured skin. This is contrary to outer edge 2, where significant reinnervation happens and total nerve fiber density at day 15 also shows non-significant change compared to uninjured skin at IE, D, and whole wound outer edge 2 (Figures 3.8 and 3.7; Table 3.1). Whereas the wound center behaves like outer edge 1 where a significant re-innervation happens by day 15, the nerve fiber density compared to uninjured skin (0.07 ± 0.007 pixels/mm^2 vs 0.13

$\pm$ 0.02 pixels/mm^2) shows significantly lower values. Stating that a considerable increase in nerve fibers at the wound outer edge 1 and wound center is still to be expected after day 15.

3.3.6 Re-innervation of the wound correlates strongly with re-epithelialization

Denervation has a detrimental effect on cutaneous wound healing. Severing the nerves hinders cutaneous wound healing, and sympathetic denervation of the skin delays re-epithelization [30]. Re-epithelization is defined by the epithelial cells migrating and growing over the wound bed and complete epithelial covering of the wound is a criterion to evaluate if a wound has healed properly [51, 69]. We investigated the correlation between nerve fiber density and re-epithelization on days 3, 7, 10, and 15 of healing and found a strong correlation ($R = 0.96$) between the two (Figure 3.9). This corroborates our hypothesis that the regeneration of nerve fibers is critical for proper wound healing in time.

3.4 Discussion

The skin is densely innervated with a complex architecture and network of cutaneous nerves, which are present in both the epidermis and the dermis. The degree of innervation has a direct effect on all the overlapping stages of wound healing [22], and previously, it has been reported that denervated wounds take a longer time to heal [70]. Therefore, precise quantification of nerve fiber density during different stages of wound healing becomes critical. PGP9.5 staining is considered as a gold standard for the quantification of nerves in skin samples in mammals [62,64]. However,

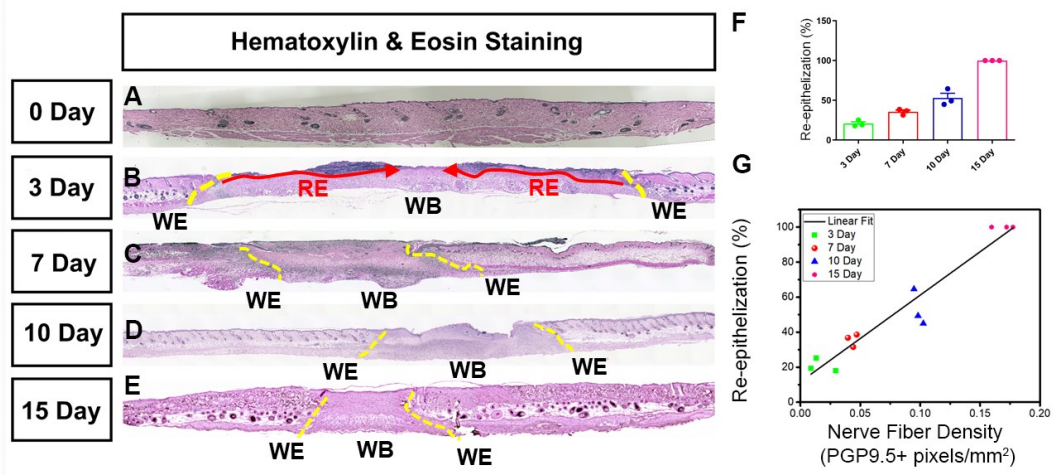

Figure 3.9: Positive correlation between re-innervation and re-epithelization. A representative image of H&E staining of the skin sample collected on (A) day 0, (B) day 3, (C) day 7, (D) day 10, and (E) day 15. The original wound edge (yellow dashed lines) on each side is determined by the absence of subdermal adipose tissue. Re-epithelialization (red arrows) is defined by epithelial cell growth. (F) quantification of re-epithelization, (G) correlation between nerve fiber density and re-epithelization at day 3, 7, 10 and 15 of wound healing. $R = 0.96$ show strong positive relation. WE wound edge, WB wound bed, RE re-epithelization.

precise quantification of cutaneous nerves is challenging because of background noise and a different pixel intensity of PGP9.5+ neurons throughout the epidermis and dermis. Therefore, to solve the problem, we used the deep neural network DnCNN for pre-processing (de-noising) of the IHC-images, followed by stringent statistical methods to formulate a threshold boundary, which is broad enough to include all the features of interest and strict enough to exclude background, so that PGP9.5+ pixels are precisely quantified in cutaneous wounds.

Applying our newly developed technique, an automated Matlab-assisted tool aided with the DnCNN, we quantified skin innervation in Female C57BL/6 mice

during normal wound healing at days 3, 7, 10 and 15. The data show that (1) the skin wound causes a substantial reduction in nerve fiber density, which significantly increased by day 15 of wound healing, (2) one of the wound edges and wound center still have significantly less nerve fiber density compared to uninjured skin at day 15 of wound healing, (3) re-epithelization and innervation share a strong correlation ($R = 0.96$).

We found that on day 3 and day 7 nerve fiber density is almost negligible when compared to the unwounded skin (Fig. 3; Table 3.1). From day 10 onwards, the nerve fiber density starts increasing marginally but reaches a significant level at the later stage of wound healing on day 15. At day 15, comparing the nerve fiber density for the whole wound bed or individually for the epidermis as well as dermis shows non-significant change when compared to the uninjured skin for the respective skin regions. Moreover, interestingly, the values are significantly higher when compared to the initial stages of wound healing, day 3, at respective skin regions (Figures 3.6 and 3.7; Table 3.1). This trend together suggests that re-innervation of the wound is initiated at the later stages of wound healing and starts becoming substantial from day 15 onwards during normal wound healing conditions in mammals. The data seems logical because for re-innervation, the proliferation of neuronal cells (Schwann cells) needs to be initiated, followed by orchestrated phenomena to develop new cutaneous nerves [71–75], and the proliferation stage of neuronal cells (Schwann cells) might overlap with the proliferation stage of wound healing, which begins approximately at day 3 and lasts for a couple of weeks. Thus, re-innervation also starts appearing substantially from day 10 onwards and shows a significant increase by day 15. Quantifying nerve fiber density separately for wound edges and wound center showed an interesting trend. At outer edge 2, similar to the whole wound, the

nerve fiber density at day 15 showed a non-significant change compared to uninjured skin at IE, D, and for the whole wound outer edge 2 (Figures 3.8 and 3.7; Table 3.1). However, for the outer edge 1 and wound center, nerve fiber density cannot reach the pre-injury level i.e., the values are still significantly less compared to uninjured skin for the whole wound (Figure 3.8). This suggests that nerve fibers are continuously innervating the cutaneous wound even beyond day 15 of wound healing. Also, on day 15 of healing, the nerve fiber density at outer edge 1 is less compared to outer edge 2 (0.12 ± 0.008 pixels/mm^2 vs 0.16 ± 0.009 pixels/mm^2) (Figure 3.8). This could be due to the difference in distance between the two outer edges of the skin wound from the cell bodies located in the dorsal root ganglia from where the cutaneous sensory nerves originate [76]. Additionally, we have been able to find a strong correlation ($R = 0.96$) between re-epithelization and nerve fiber density during time-series of wound healing, which not only corroborates the fact that the regeneration of nerve fibers is critical for proper wound healing in time but also validated our technique of using automated Matlab-assisted tool aided with DnCNN for denoising to precisely capture PGP9.5+ pixels, and thus calculate nerve fiber density.

3.5 Conclusion

This study demonstrates the use of automated deep-learning tools to accurately quantify skin innervation within the wound area, including the wound bed, wound center, and wound edges in mammals. By effectively reducing noise in immunohistochemistry (IHC) images, this technique allows for precise quantification of nerve fiber density in the epidermal and dermal layers of the wound region (bed, center, and edges). This data-centric approach provides crucial insights that can

contribute to creating predictive models for applications in precision medicine in wound healing.

The statistical approach introduced in this study serves as a valuable alternative to previously described methods [40–46]. It eliminates the need for manual data labeling, a requirement in traditional machine learning approaches for generating training data. The denoising method employed in this study has already undergone training using biological images, obviating the necessity for further training on wound-specific images. As a result, researchers can save considerable time and resources that would otherwise be expended on data collection and annotation.

The data generated shows that there is a gradual increase in nerve fiber density throughout the wound bed as well as at the wound edges, with maximum value reaching on day 15 that indicates a significant trend in re-innervation when compared to day 3 of wound healing. Additionally, on day 15 of healing nerve fiber density for the whole wound bed and at wound edge 2 reaches close to the uninjured skin. However, at wound outer edge 1 and the wound center the total nerve fiber density at day 15 of healing is still significantly less compared to the uninjured skin, asserting that re-innervation is still a continuous process beyond day 15 of healing and is important for the complete healing of the wound. The correlation between the increase in nerve fiber density and re-epithelization further supports the importance of cutaneous nerve fibers in wound healing. Overall, while our method does not elucidate the morphological characteristics of cutaneous nerves in three dimensions, it does offer a simple and cost-effective way to analyze the variations in skin innervation during different stages of wound healing.

Numerous treatments aimed at expediting and enhancing wound healing have been documented. Electric stimulation, for example, has shown promising results

in promoting innervation in human cutaneous wounds [77–79]. Many bandages e.g., Procellera claim to deliver electric fields to promote wound healing [79]. Our approach and methodology developed in this paper can assist the quantitative determination of nerve fiber density in space through the time of wound healing and facilitate assessing the effects on wound treatment. This high throughput method can be adopted for the quantification of innervation in various skin pathologies in addition to wound healing and for the quantification of innervation of other tissues and organs.

Chapter 4

Automated Neurite Identification

This chapter is from to be submitted paper title "Neural Insights: Using Deep Learning to Quantify Innervation and Assess Drug Efficacy in Wound Healing" author is a first co author of this paper.

4.1 Abstract

Wound healing is a complex physiological process encompassing hemostasis, inflammation, proliferation, and remodeling. A crucial, often underestimated, aspect of wound healing is the regeneration and innervation of nerves in the affected region. Optimal innervation is vital for reinstating normal tissue function and averting complications such as chronic pain and reduced mobility.

Quantifying the number of neurons in a wound is essential to comprehend its healing trajectory. Parameters like nerve growth, density, and neuron nucleus count are pivotal for predicting healing status. However, precise measurement of these parameters has been a formidable challenge, primarily due to the absence of an automated, accurate technique. Detecting the wound region further complicates the issue, given the variability in wound appearance depending on its origin. Traditional

methods involving manual neuron counting are not only time-intensive and reliant on expert input but can also be subjective, leading to inconsistent results.

With the advent of artificial intelligence, and more specifically deep learning algorithms, we present a novel approach to address these challenges. Our deep learning and image processing technique can automatically and efficiently detect both neurons and the wound region. By eliminating the dependence on expert judgment, our method ensures reliability, consistency, and objectivity, providing accurate neuron counts. We have validated our technique using images of neurons cultivated in cell-culture dishes. A comparative analysis with manual counting affirmed that our method surpasses the traditional technique in terms of consistency and precision.

This study seeks to utilize deep learning algorithms to underscore the pivotal role of nerve regeneration in wound healing. Our automated technique offers an expeditious, reliable, and precise means of quantifying neurons and wound regions, thereby paving the way for a more profound understanding of wound healing dynamics. This augmented comprehension is instrumental in devising superior therapeutic approaches to benefit patients.

4.2 Introduction

Wound healing is a multifaceted process orchestrated by a sequence of intricate physiological mechanisms that collaboratively reinstate tissue integrity and function. Although considerable attention has been directed towards restoring tissue integrity and functionality, the revival and reconnection of nerves within the affected area have often been overshadowed [80].

The significance of proper nerve reconnection cannot be overstated, as it

is pivotal for reinstating normal tissue function and averting complications such as persistent pain and restricted mobility. When nerve regeneration and reconnection are insufficient, patients can grapple with sensations of numbness, tingling, and other discomforts that endure long after the wound's closure. Moreover, compromised nerve renewal may lead to muscle weakness and functional deficits, curbing patients' capacity to carry out everyday activities [81].

Historically, evaluating nerve reconnection in wound healing posed challenges and necessitated invasive measures. Techniques like tissue biopsy and staining were commonplace, but their reliance on tissue removal often caused pain for patients. Nonetheless, recent strides in artificial intelligence, particularly the advent of deep learning algorithms, have unfurled novel avenues for non-invasive and streamlined assessment of nerve density and distribution [82].

Deep learning algorithms, a variant of artificial neural networks, can be educated to discern patterns within extensive datasets [83, 84]. In the context of wound healing, these algorithms can be trained on wound images to precisely gauge the concentration and arrangement of nerves in the affected zone. Scrutinizing these patterns empowers researchers to glean insights into nerve revival and reconnection during the wound healing journey [81].

Leveraging deep learning algorithms, researchers can appraise nerve reconnection in wound healing without invasive measures. This approach harbors several merits over traditional methodologies. Primarily, it sidesteps the need for tissue extraction, thus mitigating patient discomfort. Furthermore, it offers more detailed and precise measurements, as the algorithms can sift through vast data volumes and unearth nuances that might elude manual techniques [85].

Additionally, deep learning algorithms facilitate extended longitudinal inves-

tigations. By scrutinizing wound images across distinct healing stages, researchers can amass a richer comprehension of how nerve regeneration and reconnection transpire during wound healing. This methodology holds promise for unraveling nerve revival mechanisms and engendering more efficacious therapeutic strategies [86].

In essence, the quantification of nerve reconnection in wound healing bears profound implications for comprehending the holistic healing trajectory. Recent strides in artificial intelligence, particularly the evolution of deep learning algorithms, have ushered in fresh prospects for non-invasive and efficient assessment of nerve density and distribution. By harnessing these algorithms, researchers stand to uncover valuable insights into the mechanisms governing nerve revival, thereby fostering swifter and more efficient recovery. As these technological advancements continue to unfold, the horizon of wound healing appears promising, kindling optimism for enhanced recovery outcomes among patients.

4.3 Material and Methods

The primary objective of this study is to develop an automated method for identifying the area of nerve growth in images and to create a technique for assessing how nerves respond to various drugs, drug quantities, and delivery timings in laboratory experiments. To achieve this, we devised an experiment involving different drugs. In this experiment, neurons were deliberately scratched and then cultivated in a controlled environment. Images were captured on the fifth and tenth days following the scratches, and these outcomes were compared with a group of neurons that didn't receive any treatment. The central aim of this investigation is to determine the extent of innervation by analyzing these images of the scratched

neurons.

To accomplish our goal, we are employing advanced image processing methods and computer learning programs to automatically distinguish the nerve growth region. This approach is faster and more efficient compared to manual efforts. Once we have the segmented images, we will focus on the critical aspects and apply computer algorithms to quantify the extent of innervation under different conditions, such as varying drugs, drug quantities, and delivery timings.

The selection of drugs for our experiment is based on their known impacts on nerve growth and recovery. We are exercising great caution in determining the drug doses and administration times to ensure the reliability and accuracy of our findings. By capturing images at multiple time points, we can observe how innervation changes over time and make comparisons with the untreated control group.

Ultimately, the insights gained from this study could enhance our understanding of nerve regeneration and recovery processes. The automated image analysis and innervation quantification facilitated by computer techniques allow us to efficiently process a substantial amount of data, thus expediting advancements in this field. Additionally, our research holds the potential to contribute to the development of innovative drug therapies that can improve nerve regeneration outcomes, leading to better results for patients.

4.4 Experiment

The experiments were conducted by a team of researchers from Tufts University, who artificially wounded neurons using a scratch and subsequently exposed them to various treatments, as illustrated in Figure 4.1. The primary objective of

this study was to investigate the effects of different treatments on the healing process of neurons following an injury.

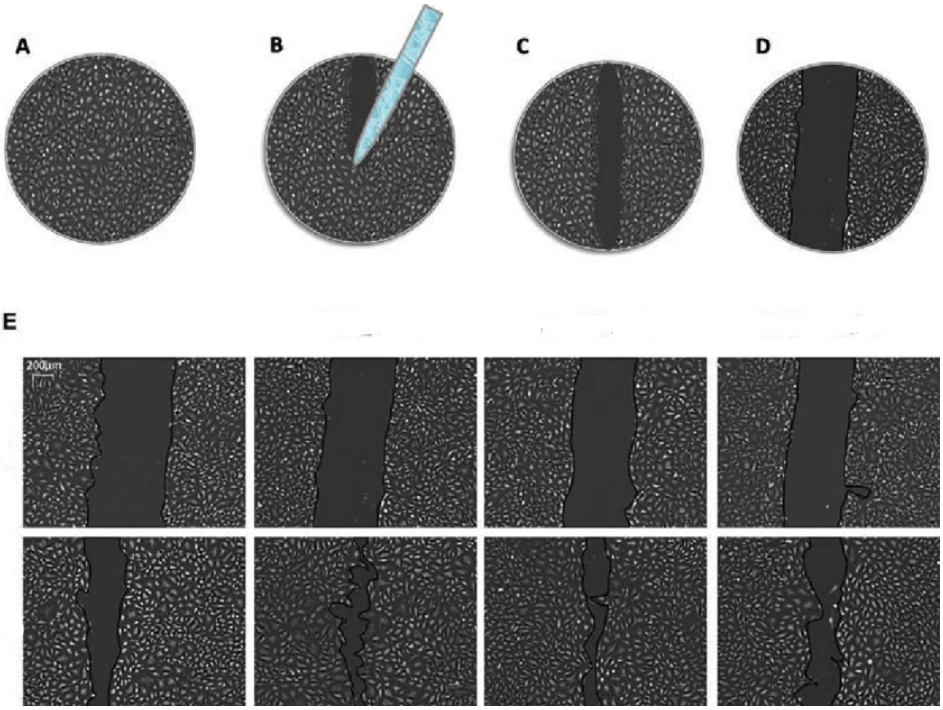

Figure 4.1: Artificially induced neuronal injury Artificial wound generation in neuron cultures

The neurons were cultured under two distinct conditions: the control group was maintained in normal culture media, while the treatment group was subjected to specific compounds or agents to observe their effects on the healing process of the wounded neurons. These treatments were administered at varying dosages and time intervals, with the goal of identifying the most effective strategy for promoting nerve growth. The team provided a set of high-quality images from the experiments to facilitate further analysis.

To accurately quantify the degree of innervation achieved by each treatment, an innovative image-based algorithm was developed. This algorithm employed advanced image processing techniques and machine learning algorithms to identify

and segment the regions of nerve growth within the images. Important features such as cell density, neurite length, and branching patterns were extracted through various methods. Machine learning algorithms were then trained using these features to accurately classify the degree of innervation associated with each treatment.

This novel image-centered approach offers significant advantages over traditional methods for measuring nerve growth. It not only enables precise and efficient quantification of innervation levels but also lends itself to easy scalability for handling larger datasets. Consequently, this methodology provides a more comprehensive understanding of how different drugs and treatments impact the process of nerve growth and regeneration.

4.5 Implementation of the Scratch Assay

The scratch assay is a unique method that combines fundamental biology with modern technology to study neuron growth dynamics and changes in their density. Such insights are invaluable for understanding how neurons regenerate and can set the foundation for potential treatments in cases of neural injuries.

The procedure begins with carefully growing and maturing human-induced neural stem cells (hiNSCs). We chose hiNSCs because of their versatility in replicating various neuronal systems. It's crucial to ensure these cells mature correctly, setting the stage for the subsequent phases of the experiment.

A key part of this assay is the custom devices developed by the Rolandi-UCSC team. These devices accurately deliver specific ions and compounds to the neuron cultures. With this precision, researchers can directly influence neuron behavior, allowing for more specific and targeted observations.

The crux of the assay involves introducing deliberate scratches or wounds on the neuron-covered surface. This simulates real-world neural injuries, making the model more relevant and comparable to actual neural injury responses.

After this injury, the neuron cultures are exposed to a range of biophysical and biochemical agents, ranging from vital ions to potential therapeutic compounds. Each is delivered under controlled conditions. The goal is to see how these agents affect neuron growth and density following the injury.

One of our main tools for observation is time-lapse fluorescence microscopy. This advanced imaging captures the ongoing responses of the neurons. With it, we can see in detail how the neurons interact with the introduced agents, giving us a clear view of regenerative processes in real-time.

In summary, the scratch assay is a comprehensive approach that merges basic biological principles with the latest technology. By using hiNSCs, precision devices, controlled injury models, and detailed observations, we gain deep insights into how neurons heal. This knowledge is crucial for developing strategies to improve recovery in cases of neural injuries.

4.6 Drug Stimulation

The study investigated the impact of zolmitriptan and serotonin, two distinct drugs, on the growth of artificially wounded neurons [87, 88]. Alongside, a control group that underwent no drug treatment and was cultivated in standard media was included. Subsequently, the neurons were subjected to imaging on the fifth and tenth days of the experiment.

To delve into the influence of varying dosages, serotonin was administered at

concentrations of 500 and 1000 micromolar, while zolmitriptan was introduced at 50 and 1000 micromolar levels. Additionally, the timing of drug administration exhibited diversity, with specific wells being treated prior to scratch induction, and others receiving treatment post-scratch formation. The obtained images were furnished by Prof. M. Levin's Lab and some examples are depicted in Figure 4.2.

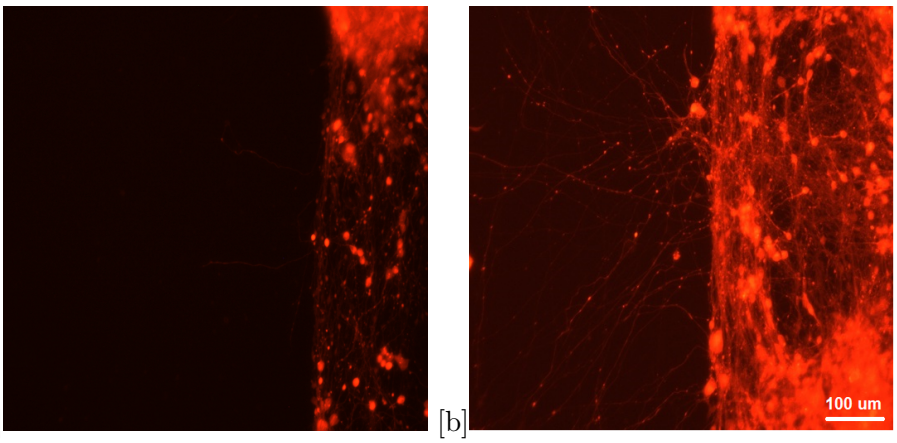

Figure 4.2: Neuron scratch in vitro. Control well with no treatment. Images captured on day 5 (a) and day 10 (b) post-scratch. Scale bar is 100 µm.

The visual data chart the progression of neuron growth over time, with images from the tenth day displaying a more pronounced surge in growth compared to those on the fifth day. Notably, the images also illustrate the presence of neuron branches, signifying that the administered drugs have spurred the generation of fresh neurites and synaptic connections.

4.7 Deep Learning Growth Area Detection

Accurately quantifying nerve growth and innervation depicted in images plays a pivotal role in comprehending the impacts of diverse drugs on neuron regeneration. This necessitates precise identification of the scratch area within each

image. To address this, a novel deep learning model was crafted to autonomously detect the scratch area in the images.

4.7.1 Labeling

Labeling the images was a meticulous manual process to ensure the utmost accuracy in our dataset. Each image was first visually assessed to get a clear understanding of potential areas showing nerve growth. With the help of Matlab, we carefully selected the regions that displayed signs of nerve tissue growth or regeneration. After identifying these areas, they were labeled as "nerve growth" to set them apart from the other regions in the image. This labeling plays a crucial role, acting as a beacon for our model to recognize similar patterns in other images. To maintain the integrity of our data, a verification step followed the labeling process. This was to confirm that the marked regions truly represent nerve growth and to ensure that no regions were either overlooked or inaccurately labeled. This detailed manual method is particularly vital in fields like ours, where precision can greatly influence the outcomes and interpretations of our research Figure 4.3.

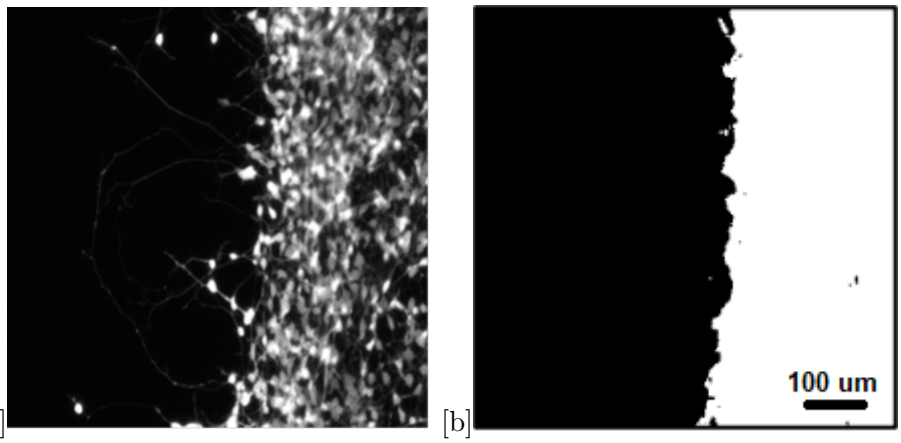

Figure 4.3: Neuron image analysis. (a) shows a black-and-white image of neurons and (b) presents the results of scratch area detection using deep learning. The scratch region is indicated in black. Scale bar is 100 µm.

The black region in Figure 4.3. b shows where new neurons are expected to appear. To enhance the diversity and strength of our training dataset, we employed standard image manipulation techniques. These included cropping, adjusting the contrast, blurring, flipping, and introducing noise to the images. By doing this, not only did we expand the volume of our training data, but we also improved the model's resilience against various disturbances, ensuring it can handle noise and other unforeseen factors efficiently.

4.7.2 Deep Learning

In our research, we employed the U-net architecture [89] for labeling regions indicating nerve growth from scratch areas. U-net is a distinguished deep learning framework specifically tailored for image segmentation (Figure 4.4). This network is fundamentally based on the fully convolutional network, but its architecture has been thoughtfully modified and expanded. Such modifications allow U-net to work

effectively with fewer training images and yet produce highly precise segmentations. Originating from the biomedical imaging domain, its design facilitates pixel-wise classification, crucial for tasks like nerve growth region detection. The architecture's strength lies in its symmetrical contracting and expanding pathways, enabling it to extract intricate details at multiple resolutions without requiring repeated iterations. Moreover, U-net's inherent design features enable it to train effectively, especially when dealing with a limited dataset, addressing a common challenge in medical imaging where obtaining abundant labeled data is often challenging.

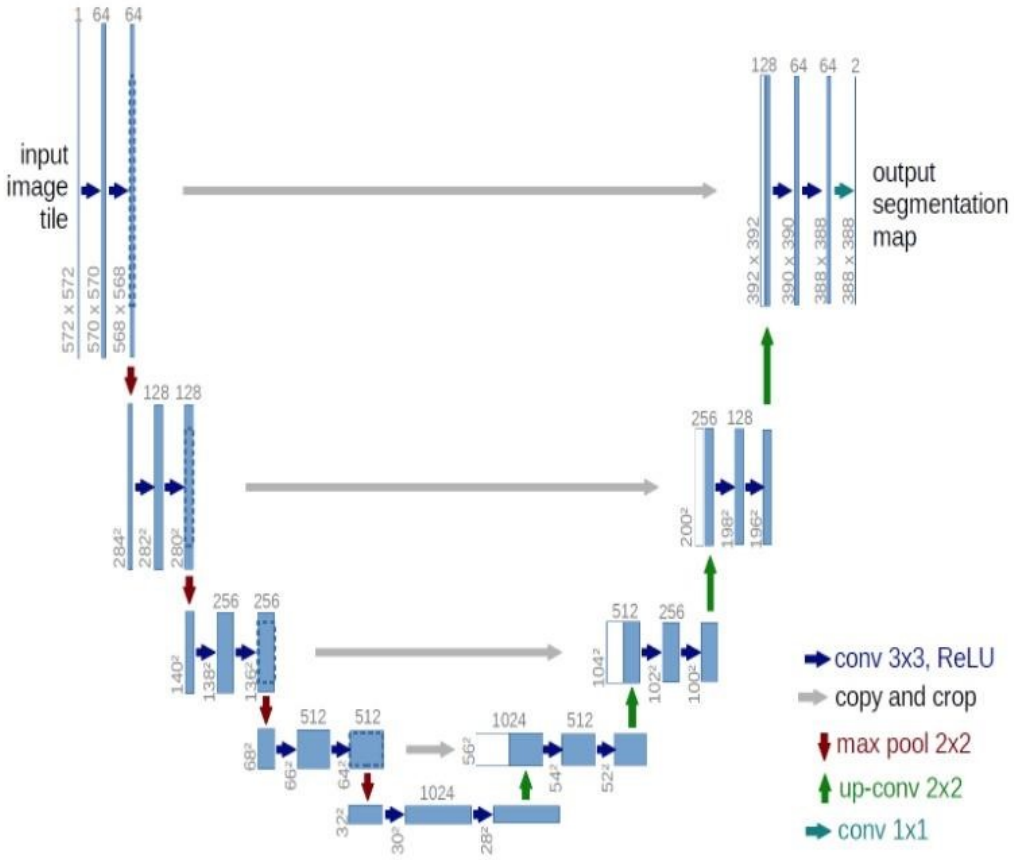

Figure 4.4: Visual Depiction of the U-net Segmentation Model.

For this study, we utilized a dataset comprising 1951 images. These images

were derived through the application of image augmentation techniques to the original dataset furnished by Prof. M. Levin's Lab. Additionally, the labeled images from the output of Matlab have been instrumental in training the U-net model. Given its tailored architecture and our specific training data, the U-net demonstrated a commendable ability to learn and discern the intricacies of nerve growth efficiently. It's worth highlighting that the efficiency of U-net is particularly pronounced when working with limited dataset images.

For the deep learning model, we utilized a standard approach to dataset partitioning to ensure efficient training and evaluation. The available images were divided into two distinct categories: training and validation images and test images. A significant majority of the dataset was allocated for training and validation. The training images form the backbone of the model's learning process, while a subset of this allocation serves for validation. Validation is essential as it continually assesses the model's performance during training, helping to fine-tune it and prevent overfitting. The remaining portion of the dataset was reserved for testing. Test images play a vital role in gauging the model's real-world performance since the model isn't exposed to them during the learning phase. This segregation ensures a well-trained model that is both validated for optimal performance and tested for its accuracy on new, unseen data.

4.8 Deep Learning Nerve Growth Detection

To study how neural growth rates change with different treatments, we needed to count the pixels that show neurons in the scratch area. At first, we used a deep learning model to tell the difference between the scratch and non-scratch parts.

Based on the results from this model, we found the growth area. In this growth area, we used MATLAB tools to spot the new neurons.

To make sure our results were right, we carefully checked the output image by hand. During this check, we removed the wrong pixels that looked like noise. We also double-checked pixels that looked like neurons, making changes where needed to make sure our nerve growth detection was accurate.

After all this work, we could see and label the pixels that showed neural growth in the scratch area. We then built a better deep learning model using the U-net structure. This new model was trained to find neurons in the scratch area, so it shows the nerves that grew there (Figure 4.5).

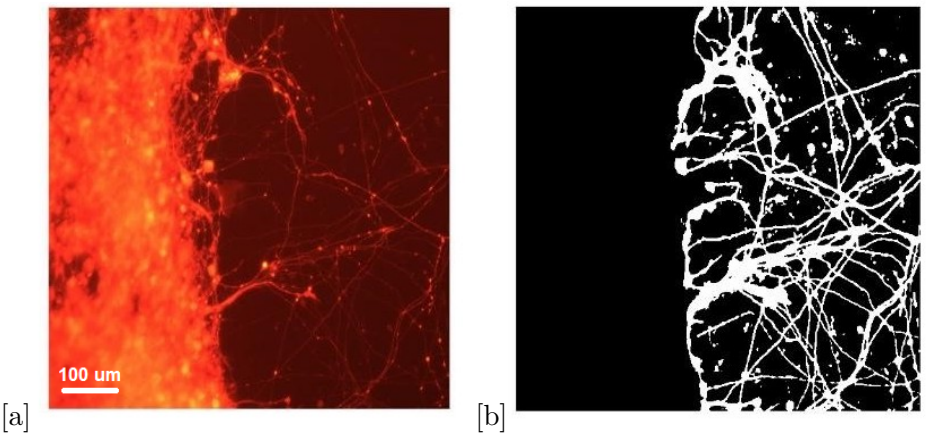

[a] [b]

Figure 4.5: (a) Side-by-Side Comparison of the Neuron Image, and (b) Highlighted Neurons Grown within the Scratch Area deep learning output. Scale bar is 100 µm.

In Figure 4.5. b, you can see the results: the white pixels are the neurons that grew in the scratch area.

4.9 Standardizing Analysis Techniques for Neuronal Growth and Innervation in Scratch Assays

In our previous section, we utilized a deep learning model to analyze in vitro responses to various drug treatments, concentrations, and delivery timings. While neurite growth is a critical parameter for innervation, it is not the sole metric. Every experimental image showcases two distinct regions: the scratch area and the intact cell culture area. It's essential to note that the size of the scratch area fluctuates across different images. Consequently, an image showcasing a more expansive scratch area might be mistakenly interpreted as manifesting more significant neurite growth. To neutralize any potential bias arising from this variability, we've anchored our quantification of innervation and cell body count to a standardized region size located at the scratch's boundary.

In our endeavor to compute neuron growth density, it's pivotal to ensure consistency in the scratch region area utilized for determining innervation density. To achieve this, we meticulously examined all the images in our dataset, searching for the smallest consistent scratch region that could be uniformly applied across images. Within this specific collection of images, the narrowest consistently replicable width was determined to be 68 pixels. Intriguingly, while the width was standardized, the vertical dimension of this region always matched the full height of the respective image, as illustrated in Figure 4.6. It's important to highlight that although the contour or edge shape of the scratch may vary across different images, the designated area for analysis remains constant.

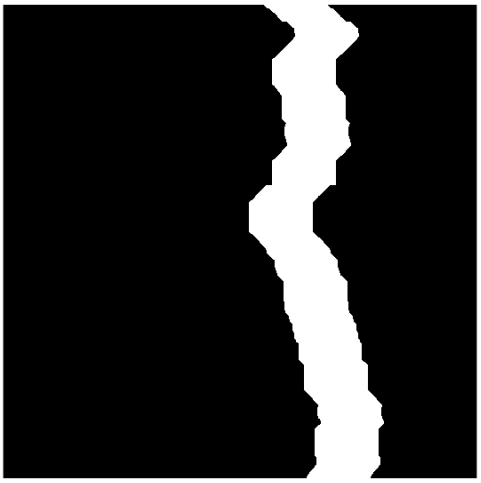

Figure 4.6: Minimum area selection in scratch assays, only part of the scratch region (68 pixels wide) is taken for further analysis.

4.10 Quantification of Neuronal Cell Bodies within the Scratch Region

Counting cell bodies provides a direct way to evaluate potential growth within the wound area. Through this method, we aim to determine whether the quantity of cell bodies serves as a reliable indicator of neuronal growth. To achieve this, we made use of images from different channels. The "Red Dye Image" presents us with a broad view of the entire neuron structure, capturing the complete context. In contrast, the "Blue Dye Image" has been designed specifically to emphasize and bring out the cell bodies, ensuring they are distinct and countable. For a hands-on understanding and visualization of how these images facilitate the identification and enumeration of cell bodies, Figure4.7 offer a clear depiction of the cell bodies prominently showcased.

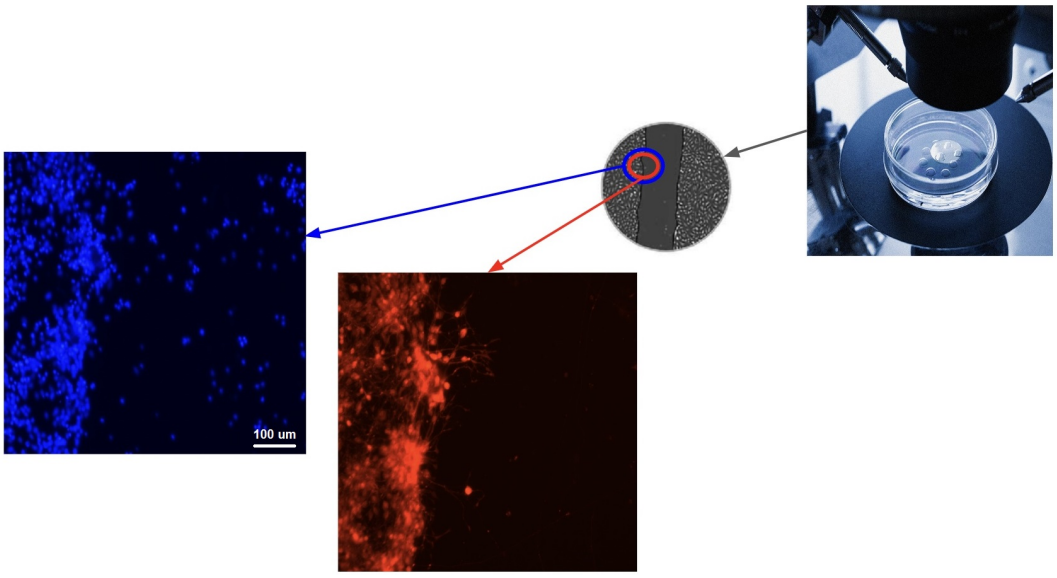

Figure 4.7: Compilation of Neuronal Images: Cell Bodies (Blue) and Neuron Nucleus and Branches (Red).

4.10.1 Cell Body Detection Method

To ascertain the number of cells within specific regions of our neuron images, we devised a sophisticated image processing algorithm. This algorithm hinged on the capabilities of several Python libraries: Trackpy [90]: A high-dimensional particle-tracking package. Matplotlib [91]: Renowned for its visualization capabilities. Numpy [92]: Facilitates diverse mathematical operations on arrays. Our primary task was to represent the image as an n-dimensional array, detailing intensity values ranging between [0, 255] for an 8-bit image. Once the image was processed, our focus shifted to identifying and cataloging the salient features of each cell, such as its spatial positioning and appearance traits like mass and size, denoted in pixels. To fine-tune our detection capabilities and minimize errors, we employed filters based on mass and size.

The algorithm followed these core steps: Grayscale Conversion: The image

was transitioned to grayscale to simplify intensity analysis. Binarization: A threshold was set, beyond which the image underwent binarization. Connected Component Analysis: This technique facilitated the recognition of individual cells. Cell Feature Calculation: Each identified cell's mass and size were computed. Filter Application: To ensure precision in cell detection, we applied filters based on the computed mass and size, eliminating abnormally small or large cells. Cell Counting: Finally, cells within our defined region of interest were counted. Upon implementation in Python, the algorithm was subjected to rigorous testing using a diverse set of neuron images. In every instance, it demonstrated remarkable accuracy in identifying the correct number of cells within the defined regions.

4.10.2 Detecting the Number of Cell Bodies

To validate our algorithm's effectiveness, we began by determining the total count of cell bodies within the scratch area. Subsequently, we'll be contrasting this count with the cell body density found in a predetermined fixed area.

For accurate cell counting, we employed an algorithm which had been previously designed for macrophage image processing.

4.10.3 Determination of Cell Body Density in Neuronal Scratch Assays

In our endeavor to produce consistent and reliable data, we employed standardized analysis techniques for evaluating neuronal growth and innervation in scratch assays. Central to this approach is the concept of cell body density. Rather than merely counting the absolute number of cell bodies, we sought to understand their concentration in a specific area, providing a more nuanced view of neuronal

response.

To achieve this, we focused on the minimum scratch edge area we had previously identified. Within this confined region, we undertook a methodical count of the cell bodies. This count, when normalized to the area's size, can offer us better insight into the connection between the cell body density and growth. By using such a standardized method, we ensure both the reproducibility of our results and their alignment with established analytical benchmarks in the field.

4.11 Results

In our pursuit to understand the intricate response of neurons to various treatments, we delved into quantitative analyses using advanced pixel detection methodologies. This enabled us to assess the effect of specific drugs on neuronal growth, thus paving the way for potential therapeutic applications. Within this framework, we examined two primary compounds: serotonin and zolmitriptan. Their impacts on neuronal growth were analyzed under varied conditions and dosages, and the outcomes of these investigations are presented below.

4.11.1 Serotonin Treatment

In our first experiment, neurons were pre-treated with a 500 µM concentration of serotonin 24 hours prior to the scratch. This was complemented by an additional serotonin dosage on the day the scratch was made. As we evaluated the outcomes depicted in Figure 4.8, we noted the variance in drug concentrations in the medium post-scratch.

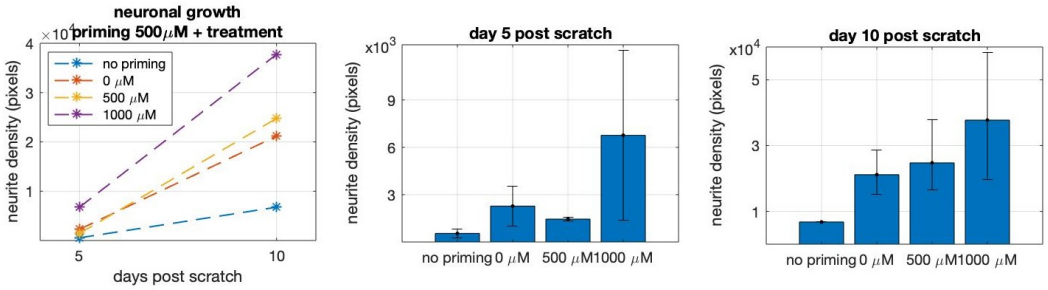

Figure 4.8: Innervation Analysis in Scratch Assay: Effect of Serotonin Dosages Post 500 µM Pre-treatment, Highlighted with 10 and 90 Percentile Error Bars.

From our assessments with serotonin, it emerged that a concentration of 1000 µM had the most pronounced impact on promoting neuronal growth. In our subsequent experiment focused on serotonin treatment, neurons received the treatment just post-scratch, with variations on the day of application. Observations from Figure 4.9 provide insights into how treatment day and dosage influence innervation.

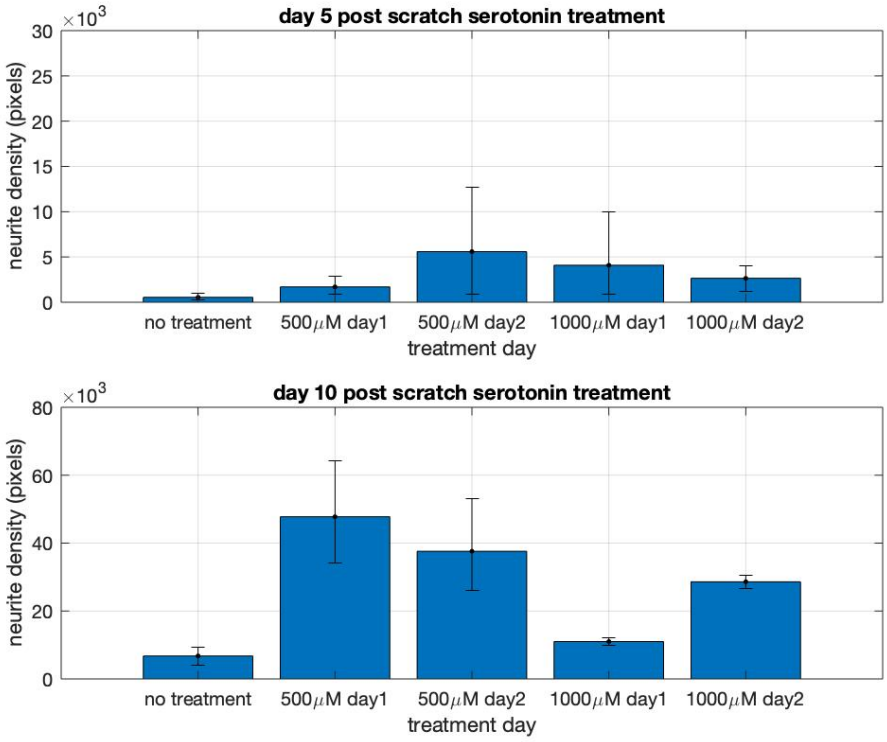

Figure 4.9: Deep Learning Analysis of Scratch Innervation: Varying Dosages and Application Days of Serotonin Post-Scratch. Error bars represent 10 and 90 percentiles.

Notably, the scratch treated with 500 micromolar serotonin exhibited maximum growth by the 10th day. However, it's worth noting that certain images presented significant discrepancies, prompting the need for further validation of some results.

4.11.2 Zolmitriptan Treatment

For the Zolmitriptan treatment, a two-step process was employed. Initially, a priming concentration of 50 µM of Zolmitriptan was introduced 24 hours prior to creating the scratch. This served as the pre-treatment phase. Subsequent to the scratch, the media was updated to include varying concentrations of Zolmitriptan to

investigate its impact. The treatment timeline corresponding to this is illustrated in Figure **??**. An analysis of the Zolmitriptan pre-treatment experiments is shown in Figure 4.10.

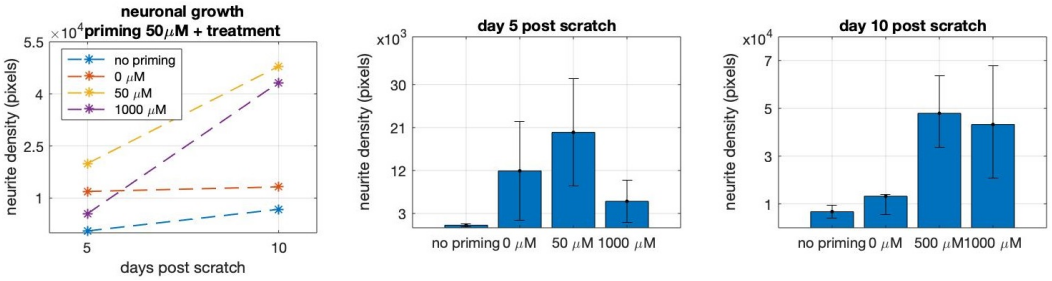

Figure 4.10: Neuronal Growth Analysis Post Zolmitriptan Treatment: Deep Learning Quantification Following 50 µM Priming. Error Bars Indicate 10 and 90 Percentiles.

Figure 4.10 reveals that the 50 µM concentration yielded the maximum neuronal growth.

For the subsequent study with zolmitriptan, the drug was administered post-scratch, with variations in both dosage and timing across different samples. The regimen detailing dosage schedules and imaging intervals can be seen in Figure **??**. As depicted in Figure 4.11, the highest neuronal growth was observed on the 10th day post-scratch for samples treated with 50 micromolars of zolmitriptan.

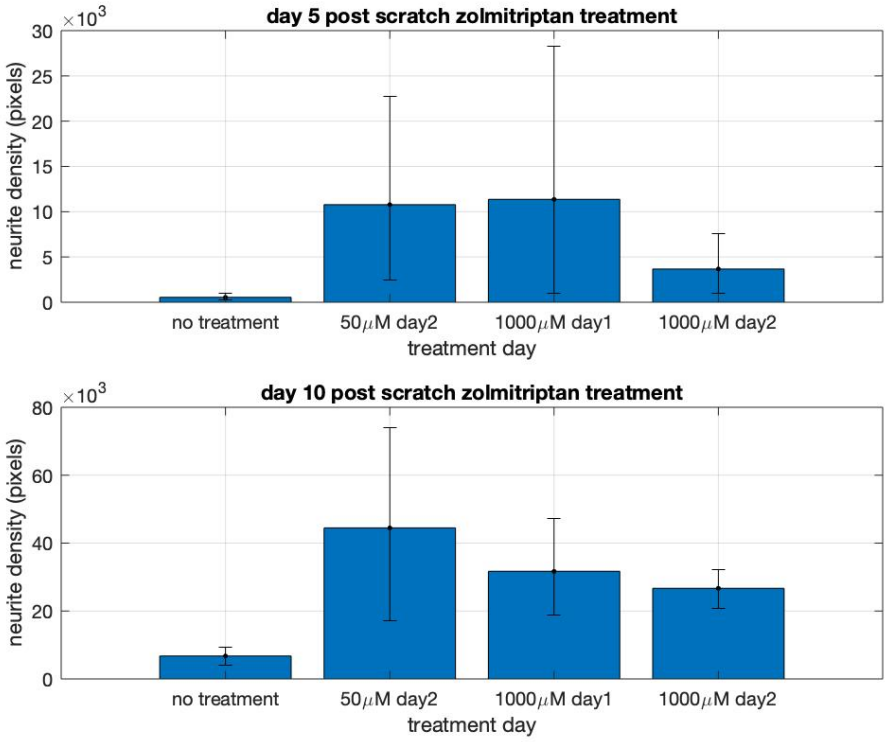

Figure 4.11: Innervation Quantification via Deep Learning Analysis: Variations in Neuronal Growth with Diverse Dosages and Timing of Zolmitriptan Post-Scratch Treatment. Error bars represent the 10th and 90th percentiles.

Intriguingly, as the zolmitriptan dosage was ramped up, there was a noticeable decline in growth. However, even at the heightened dose of 1000 micromolar, the growth outpaced that of the control well, which had no zolmitriptan exposure.

4.11.3 Standardizing Innervation Density

To attain a precise understanding of drug and dosage effects on innervation, we've standardized our analysis by focusing on the minimum scratch area. By employing a deep learning model, we have quantified the mean neurite density across the image collection. Additionally, the 10th and 90th percentiles provide further

insights into the range and distribution of results.

To identify the most effective drug, we compared the effects of zolmitriptan and serotonin on innervation. displays the neurite density observed 5 days post-scratch following treatment with either serotonin or zolmitriptan. Figure 4.12 presents the outcomes on day 10. The data in Figure 4.13 reveals that zolmitriptan at a concentration of 1000 µM (mean value) induces a greater neurite outgrowth compared to both serotonin and the control.

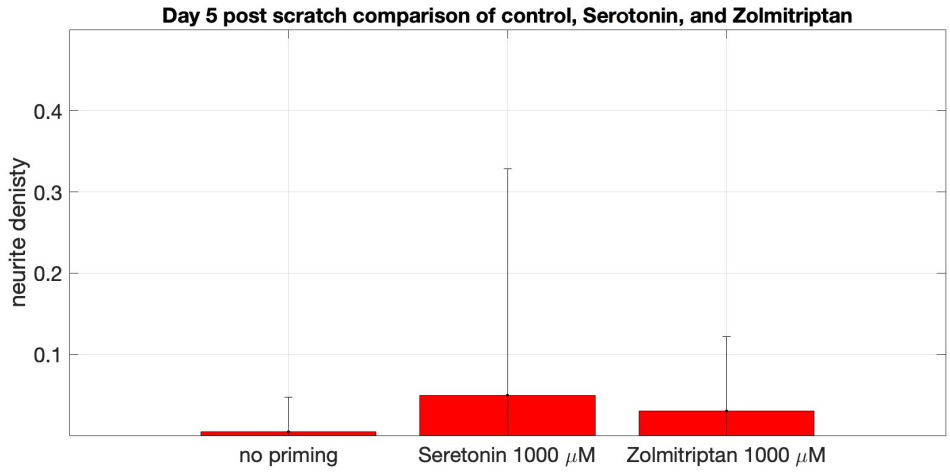

Figure 4.12: Day 5 Neurite Density Post-Scratch: Comparing Control, Serotonin, and Zolmitriptan Treatments. Error Bars Indicate the 10th and 90th Percentiles.

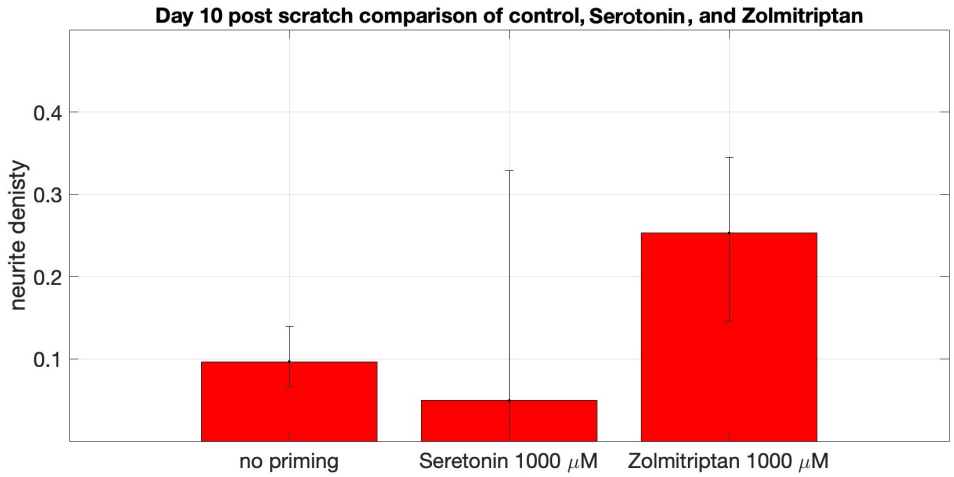

Figure 4.13: Innervation Density Comparison Post-Scratch: Control vs. Serotonin vs. Zolmitriptan. Error Bars Represent the 10th and 90th Percentiles.

Furthermore, this significant increase in neurite outgrowth underscores the potential therapeutic advantages of zolmitriptan over serotonin, warranting further exploration into its mechanisms and broader implications in neuroregenerative medicine.

4.11.4 Quantification of Cell Body Numbers

Figures 4.14 and 4.15 provide a comparative analysis of cell body counts between serotonin and zolmitriptan experiments.

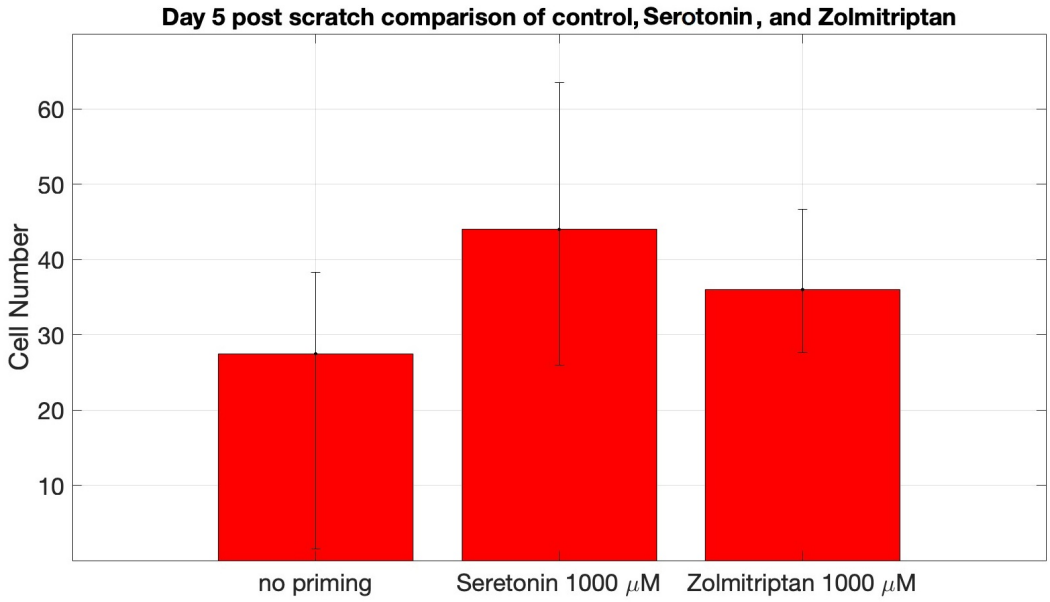

Figure 4.14: "Day 5 Post-Scratch Cell Counts: Comparisons Among Control, Serotonin, and Zolmitriptan Treatments. Error Bars Represent the 10th and 90th Percentiles from Replicates.

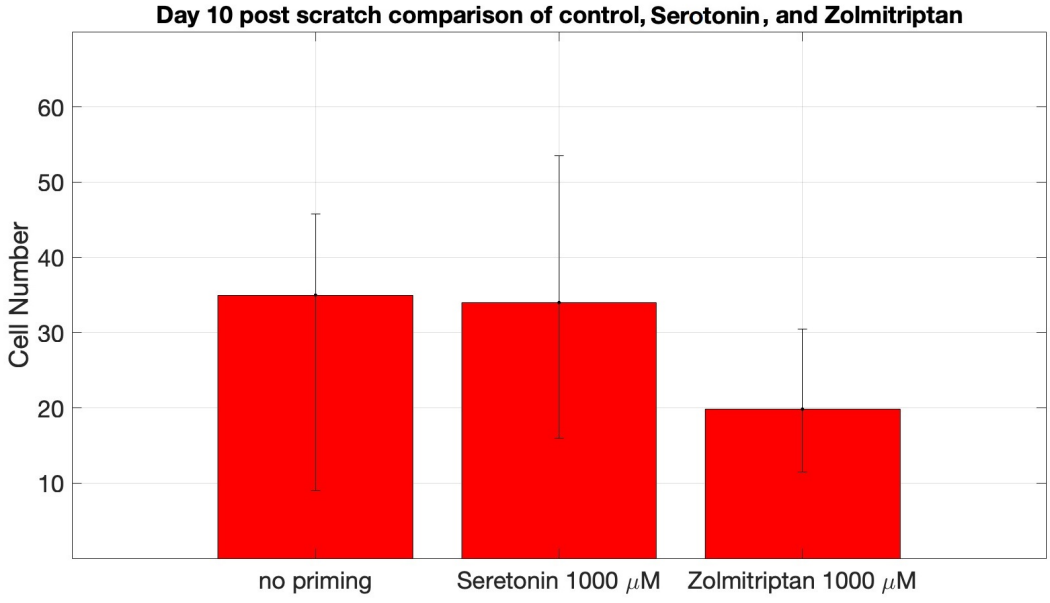

Figure 4.15: "Day 10 Cell Counts Post-Scratch: A Comparison of Control, Serotonin, and Zolmitriptan Treatments. Error Bars Denote the 10th and 90th Percentiles from Replicates.

The data underscores the relative consistency in cell body counts across both treatments at different time points post-scratch.

4.11.5 Quantification of Cell Body Density Numbers

To calculate the cell body density, we tallied the number of cell bodies within the specified minimum scratch edge area, same area that we used in the Neurons growth minimum area.

When contrasting the effects of serotonin and zolmitriptan, it's evident that the condition without priming yields the highest average cell count within the minimum scratch area on day 10. Yet, on day 5, the cell counts appear to be similar, as illustrated in Figures 4.16 and 4.17.

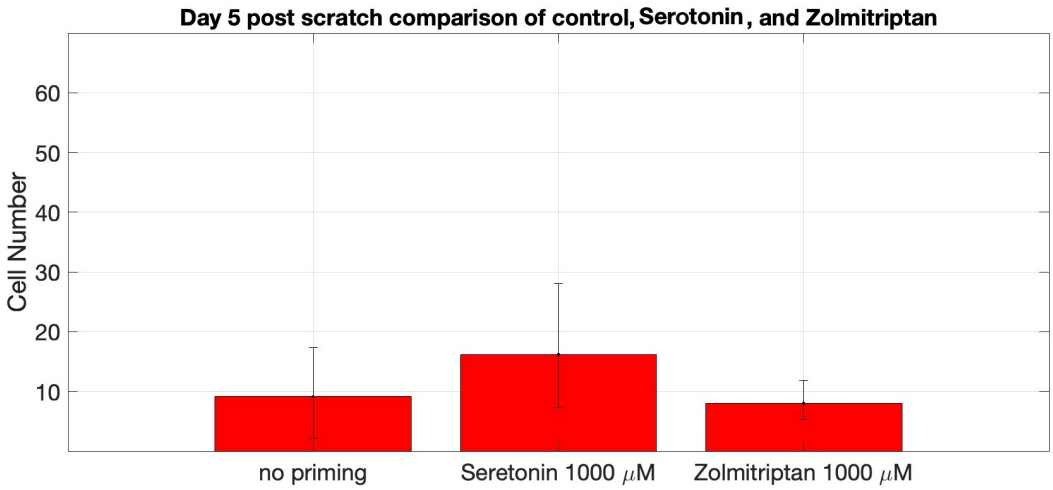

Figure 4.16: Comparative Effects of Serotonin and Zolmitriptan on Cell Body Count in the Minimum Scratch Area on Day 5.

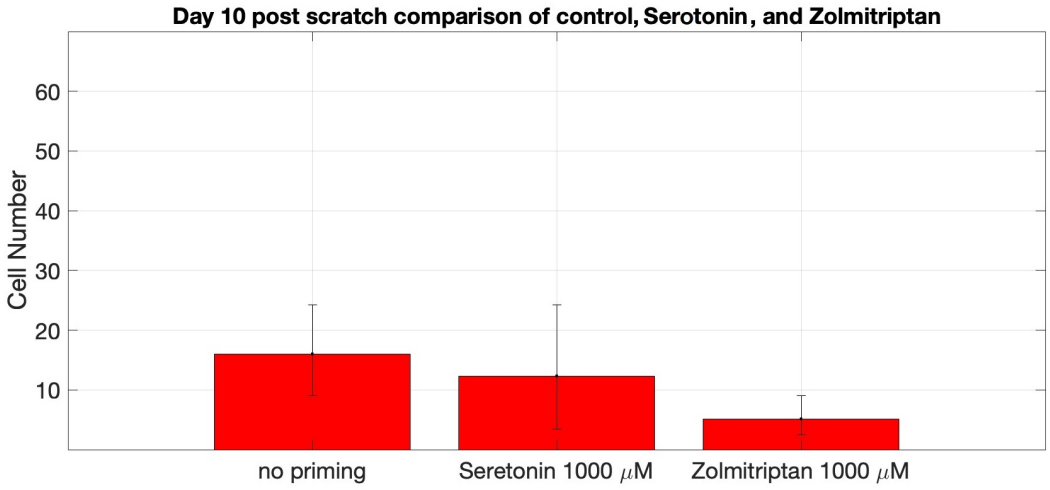

Figure 4.17: Comparative Effects of Serotonin and Zolmitriptan on Cell Body Count in the Minimum Scratch Area on Day 10.

Upon comparing the outcomes, it is evident that treated cells typically migrate further from areas of high density over time, leading to enhanced neurite outgrowth. Conversely, in the control scenario, cells predominantly remain close to the scratch boundary, resulting in subdued neurite growth.

4.12 Conclusion

The intricate process of wound healing is greatly influenced by nerve regeneration and innervation, which, if not properly assessed, can hinder our understanding of wound healing dynamics and the development of effective therapeutic solutions. A major challenge in this realm has been the quantification of neurons in wounds due to the lack of automated techniques and inconsistencies in manual methods.

In our study, we employed deep learning algorithms and image processing techniques to establish a new standard in the quantification of innervation, focusing particularly on the healing trajectory. We demonstrated that our novel method is not

only faster and more efficient but also more consistent and reliable than traditional manual neuron counting. Our method was validated using neuron images from cell-culture dishes, and comparative analyses with manual counting showcased its superior precision.

Through meticulous experimentation, our research highlighted the significant influence of drugs like serotonin and zolmitriptan on nerve regeneration. Notably, both serotonin and zolmitriptan exhibited the capability to enhance neurite outgrowth significantly. However, our findings suggest zolmitriptan's potential as a more effective therapeutic agent over serotonin in promoting neurite outgrowth, thus emphasizing the importance of further exploration into its broader applications in neuroregenerative medicine. This insight could redefine therapeutic approaches and offer substantial benefits to patients.

The data also unveiled the consistent cell body counts across both serotonin and zolmitriptan treatments. This stability in neuronal cell body count underscores that the drugs primarily influenced neurite outgrowth without significantly affecting the underlying cell numbers.

An intriguing observation was the association between increased drug concentration and decreased average cell body density in the zolmitriptan treatment, whereas serotonin treatment did not showcase a similar pattern. Further research might provide insights into the underlying mechanism responsible for this trend.

Furthermore, the differences in neurite outgrowth between treated cells and control conditions suggest that cells in treated environments likely migrate farther from dense areas, promoting enhanced neurite growth. In contrast, cells in control settings mainly remain closer to the scratch boundary, thus restricting neurite extension.

In summary, our study underscores the transformative potential of artificial intelligence in advancing our comprehension of wound healing dynamics. By leveraging the power of deep learning, we have forged a path towards more reliable and objective quantification of innervation, providing an invaluable tool for researchers and clinicians. As the field of neuroregenerative medicine evolves, these insights will undoubtedly play a pivotal role in the development of next-generation therapeutic solutions for improved wound healing.

Chapter 5

Dynamic Optimization

5.1 Abstract

In the preceding chapters, we explored the foundational concepts of personalized precision treatments and the challenges posed by patient variability. As we journey into Chapter 5, we delve deeper into cutting-edge solutions that bridge the gap between intricate patient-specific requirements and advanced control strategies. This chapter unveils transformative advancements, unveiling a holistic approach that integrates artificial intelligence, deep learning, and reinforcement learning which is a novel work in personalized medicine.

5.2 Transformative Advances in Personalized Precision Treatments

In the ever-evolving landscape of modern medicine, personalized precision treatments have become the focal point of research, propelled by the remarkable progress in artificial intelligence [93–95]. The urgency for precision treatment arises

from the inherent variability in patients' responses to medications, reflecting molecular disparities both among individuals and within the same person across different contexts [96]. Personalized treatments strive to decipher bespoke dosages of drugs, appropriate drug types, and optimal drug delivery timings for each patient. This tailoring is meticulously crafted based on current and predicted patient responses, weaving a rich tapestry of experimental data and rigorous statistical analysis into the fabric of medical decision-making [97].

However, amidst the complexities of personalized precision treatments lies a significant challenge: the design of optimal treatments that account for multifaceted factors such as gender, age, weight, and the intricate interplay of pharmacokinetic and pharmacodynamic elements [98]. Traditional controller design methods often stumble in the face of such complexities. This challenge finds innovative solutions in the realm of reinforcement learning (RL)-based approaches, offering real-time control algorithms capable of optimizing and adapting drug dosing amidst the intricate maze of patient variability [99].

5.3 Introduction

Deep Reinforcement Learning (DRL) has emerged as a transformative approach in various fields, including healthcare and biological systems. In the specific context of wound healing, DRL offers a novel and promising avenue for optimizing treatment strategies. The crux of implementing DRL in wound healing lies in its ability to interact dynamically with the biological system, adapting to real-time data and iteratively improving its decision-making process.

The controller, a central component of the DRL framework, is responsible

for gathering and interpreting data from the biological system. It closely monitors the healing process, constantly comparing the observed responses of the wound with a set of predefined, desirable healing trajectories. This comparison is crucial as it forms the basis for the controller to make informed decisions about the next course of action.

However, the challenge in this application is not just in the collection and interpretation of data, but also in establishing a method to accurately compare healing timelines. This involves developing a robust model that can predict wound stage probability profiles, which are then matched against real-world data. The aim is to find the best correspondence between the predicted stages of wound healing, as per the model, and the actual progression observed in real cases.

By leveraging DRL, the system can not only predict but also adapt its strategies to align with the optimal healing path. It learns from each interaction, adjusting its parameters and strategies based on the success or failure of previous actions. This adaptive learning mechanism is key to handling the complexities and variabilities inherent in biological systems, particularly in the context of wound healing where each case can present unique challenges.

5.4 The Mathematical Model

To provide an in-depth understanding of the mathematical model being utilized for training Deep Reinforcement Learning (DRL) in the context of wound healing, it's important to delve into the core aspects of the model and its relevance to the biological process it represents. The mathematical model, in this scenario, acts as a crucial bridge between theoretical understanding and practical application,

especially in the nuanced field of wound healing.

5.4.1 Ordinary Differential Equation Model

The mathematical model at the heart of this application is designed to simulate the complex dynamics of wound healing. Wound healing is a multifaceted biological process involving several stages: hemostasis, inflammation, proliferation, and remodeling. Each of these stages is governed by intricate biological processes and interactions among various cells, growth factors, and extracellular matrix components. The mathematical model encapsulates these stages through a set of equations and parameters that reflect the underlying biological mechanisms [100]. For instance, differential equations can represent the rate of cell proliferation, migration, and the formation of new tissues. Parameters within these equations are adjusted to reflect the variations in healing rates due to factors like age, health conditions, and the severity of the wound.

The core of the model consists of differential equations that describe the behavior and interaction of various biological factors. These equations can model the proliferation and migration of cells, the formation and degradation of extracellular matrix, and the action of growth factors and enzymes. For example, the rate of fibroblast proliferation might be represented by a logistic growth equation, which factors in the limited resources and space available in the wound area.

Wound healing is both a spatial and temporal process. The model must account for changes over time as well as spatial variations within the wound. Partial differential equations are often used to incorporate these aspects, allowing the model to simulate how cells and substances move and change throughout the wound area over time.

Parameters in the equations are critical as they represent the rate constants, diffusion coefficients, and other biological characteristics pertinent to wound healing. These parameters are typically derived from experimental data or literature. They are crucial for the model to accurately reflect the biological realities of wound healing. The equation for the wound healing model has been provided based on the [100].

$$\dot{D}_b = -k_1 D_b M_1 \tag{5.1}$$

$$\dot{M}_1 = k_2 D_b - k_1 D_b M_1 - \frac{k_4 M_1^q}{K^q + M_1^q} - k_{d1} M_1 + D_1 \left(\frac{1}{r} \frac{\partial M_1}{\partial r} + \frac{\partial^2 M_1}{\partial r^2} \right) \tag{5.2}$$

$$\dot{M}_2 = \frac{k_4 M_1^q}{K^q + M_1^q} - k_{d2} M_2 + D_1 \left(\frac{1}{r} \frac{\partial M_2}{\partial r} + \frac{\partial^2 M_2}{\partial r^2} \right) \tag{5.3}$$

$$\dot{C} = k_5 M_2 - k_r C \tag{5.4}$$

$$\dot{N} = C\alpha N(1 - N) + D_n \left(\frac{1}{r} \frac{\partial N}{\partial r} + \frac{\partial^2 N}{\partial r^2} \right) \tag{5.5}$$

In these equations:

- $\dot{D}_b, \dot{M}_1, \dot{M}_2, \dot{C}, \dot{N}$ represent the rates of change of wound debris, M1 macrophages, M2 macrophages, temporary tissue, and new tissue, respectively.
- The constants $k_1, k_2, k_4, k_{d1}, k_{d2}, k_5, k_r, \alpha, D_1$, and D_n are each associated with specific processes in wound healing dynamics:
  - k_1 with the elimination of debris ($k_1 DbM_1$),
  - k_2 with the attraction of M1 macrophages by debris,

  - k_4 and k_{d1} with the polarization and natural death of M1 macrophages, respectively,
  - k_{d2} with the death of M2 macrophages,
  - k_5 with the production of temporary tissue by M2 macrophages,
  - k_r with the destruction of temporary tissue,
  - α with the growth rate of new tissue,
  - D_1 with the spatial migration of M1 macrophages,
  - D_n with the migration rate of new tissue cells.

- M_1, M_2, D_b, C, N denote the concentrations or populations of M1 macrophages, M2 macrophages, wound debris, temporary tissue, and new tissue.
- The terms involving partial derivatives $\frac{\partial}{\partial r}$ and $\frac{\partial^2}{\partial r^2}$ account for spatial distribution and diffusion effects in the cylindrical coordinate system.

To ensure the model's accuracy, it is calibrated against real-world data from clinical observations or experimental studies. This process involves adjusting the parameters until the model's outputs align closely with observed healing patterns. This validation step is essential for ensuring that the model is a reliable representation of the actual healing process.

The parameters chosen for solving this ordinary differential equation (ODE) model are detailed in table5.1.

5.5 Integrating the ordinary differential equation model

The ordinary differential equation (ODE) model serves as a quantitative tool to simulate the dynamic interactions between different cells and tissues involved

Table 5.1: The values of parameters used in experiments

Parameter	R	L	T	β	ρ	κ	q	γ_1	γ_2	μ	D	$\tilde{D}_n$	$\tilde{\alpha}$	γ	$\epsilon_{\min}$	σ	ν
Value	3 mm	0.03 mm	1/3 day	1	0.1	0.05	5	0.1	0.1	0.2	0.32	3×10^{-4}	1.8	0.995	0.01	0.95	0.99

in wound healing. It not only helps in understanding the inherent mechanisms but also provides a means to visualize how interventions might alter the healing trajectory. The primary objective of the research is to harness this model to find ways to expedite the healing process, which in turn could significantly reduce recovery time for patients.

A critical factor in enhancing wound healing is the timing of therapeutic interventions. The sensitivity of the healing speed to treatment timing necessitates a strategic approach to when drugs should be administered. To optimize this timing, comparisons of various treatment schedules are made using the ODE model, aiming to identify which yields the fastest healing.

Integrating the ODE model with real-time data from wound healing observations allows for a more informed decision-making process regarding the timing and expected outcomes of treatments. This integration is crucial for personalized medicine, where treatment dosages and schedules can be tailored to the individual's response to therapy.

To effectively halve the wound healing time, a reference timeline—representing the normal healing process without intervention—is required. This serves as a benchmark to measure the effectiveness of treatments. Furthermore, developing a reference model for the post-treatment scenario enables predictions of expected outcomes following therapeutic interventions. By comparing real-time healing data with these reference models—both with and without treatment—researchers and clinicians can evaluate whether the drug is achieving the desired effect.

The upcoming discussion in the document will delve into the construction of reference models for both treated and untreated wounds and elucidate the process of mapping real-time wound data onto these models. This is essential for assessing

the efficacy of treatments and making any necessary adjustments to the therapeutic regimen.

5.6 Standardizing Wound Healing Assessment: Utilizing Ordinary Differential Equation Model

To effectively monitor the progression of wound healing, it is crucial to establish a benchmark for comparison. The Ordinary Differential Equations (ODE) model provides such a benchmark by simulating the chronological changes in the wound environment from the moment of injury to complete healing. Specifically, the model delineates the temporal dynamics of macrophage activity—a critical cellular component in the healing process.

Our approach posits that while each patient's wound undergoes the same sequence of biological events, the rate of progression through these stages varies. To illustrate this, we refer to Figures 5.1 and 5.2 , which depict the normalized activities of macrophages and tissue components in the wound healing model over time. The X-axis represents time in days, tracking the progression from the onset of the wound healing process to 20 days and 12 days post-wound initiation. The Y-axis indicates the activity levels of each entity, normalized to their respective peak values observed during the simulation. A value of 1 on the Y-axis indicates 100 percent of the peak activity for that entity. The graph shows the dynamics of debris clearance (blue), M1 macrophage activity (red), M2 macrophage activity (yellow), temporary tissue formation (purple), and new tissue growth (green). These entities represent the biological processes occurring at the wound site, with macrophage activities reflecting their roles in inflammation and repair, and tissue lines indicating the presence and

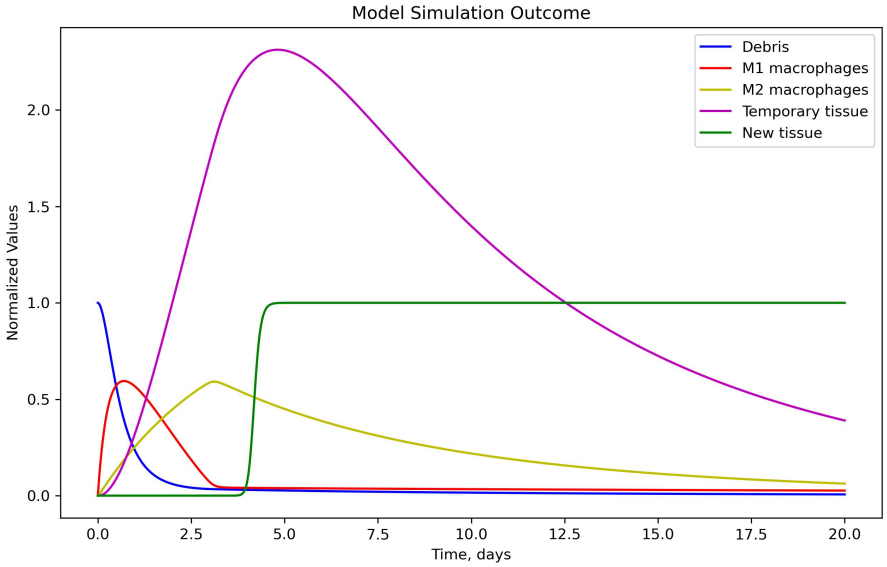

Figure 5.1: Normalized activities of macrophages and tissue components in the wound healing model over a 20-day period. The Y-axis represents the activity levels of each entity, normalized to their respective peak values observed during the simulation.

maturation of wound tissue. Notably, Figure 5.1 indicates the peak of M1 macrophage activity occurs near day 1, while 5.2 illustrates a faster wound healing model where this peak is reached at approximately half a day. Although the underlying biological dynamics remain consistent, the time scales differ significantly. By using the day as a unit of measurement instead of the specific wound stage, we establish a standardized reference point that is consistent across different patients. For instance, even though the peak of M1 macrophage activity in Figure 17 is reached more rapidly, we can standardize this as 'day 1' in the healing timeline. This standardization allows us to compare the healing progression of different wounds on a unified temporal scale. Consequently, each patient's healing rate can be assessed against this model, providing a universal metric for evaluating the efficacy of treatments and interventions.

In essence, this ODE model-based standardization affords us a method to

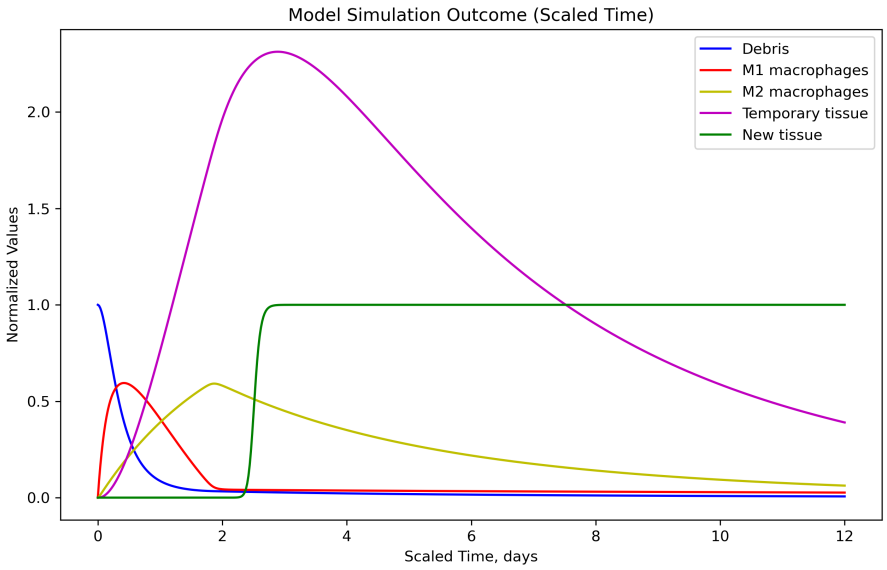

Figure 5.2: Normalized activities of macrophages and tissue components in the wound healing model with higher healing rate over 12 days. The Y-axis represents the activity levels of each entity, normalized to their respective peak values observed during the simulation.

quantify and compare the healing trajectories across diverse patient populations, offering a more nuanced understanding of individual healing rates. This understanding is pivotal for optimizing treatment protocols and tailoring them to accelerate the healing process for each patient uniquely.

Utilizing advanced machine learning techniques, specifically a non-linear Radial Basis Function Support Vector Machine (RBF SVM), we establish a correlation between macrophage concentration levels and temporal stages of wound healing. This sophisticated algorithm interprets outputs from the ODE model, assigning day-based labels that correspond to distinct stages in the wound recovery process (as illustrated in 5.3). These day-based labels serve as milestones against which the progression of wound healing can be gauged. By applying this method, we can accurately predict the day label for a given state of wound healing, enabling a comparison of observed wound stage probabilities with these predictive labels. Such a comparison is instrumental in assessing the effectiveness of healing over time, providing a quantitative measure to monitor recovery and inform treatment decisions.

When dealing with wound healing, the ODE model provides the temporal dynamics of the healing stages, but it might be non-linear and complex due to the interaction of various biological processes. An RBF-SVM can learn this complex relationship by fitting a model that can predict, for example, the stage of healing based on the current state of the wound.

Specifically, it can take the output probabilities of the ODE model for different stages of wound healing and use them as inputs to predict the time or stage of the wound. The RBF kernel allows the SVM to capture the non-linear patterns in how these probabilities change over time, providing a prediction of when a wound will reach a particular stage in the healing process, which can then be used for making

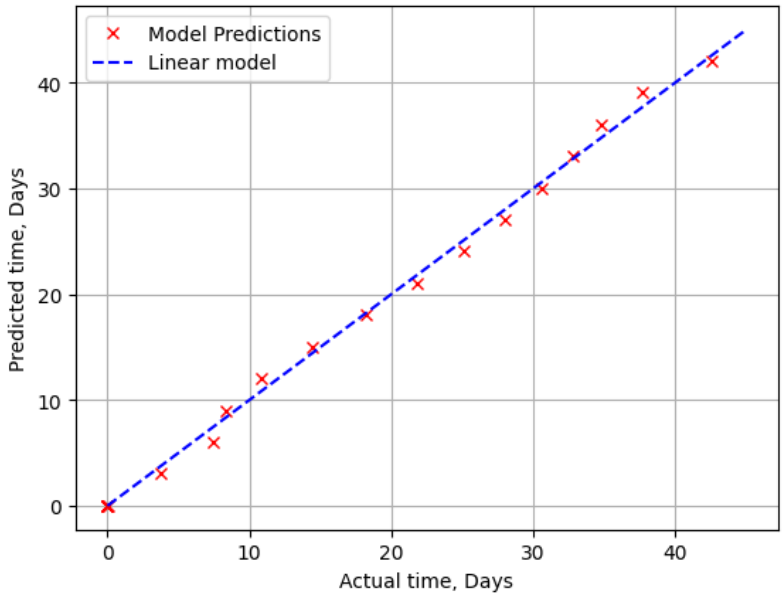

Figure 5.3: Linear model predicts wound time based on variables values. Predicted vs actual time is shown.

clinical decisions.

5.6.1 Reference model after treatment

When therapeutic interventions are introduced to the wound healing environment, they invariably alter the established timeline of cellular activities, particularly those of macrophages. These changes, as delineated in the preceding section, can accelerate the healing process. By leveraging the previously established linear model—originally fitted to the untreated ODE model ([100])—we can extend this framework to include scenarios where treatment is administered. This allows us to predict new time labels that reflect the accelerated healing trajectory post-treatment.

The treatment's effectiveness is reflected in the rate at which the predicted time labels are reached compared to the untreated model. Essentially, the treated model's time labels will be 'ahead of schedule,' indicating a faster progression through

the healing stages. By conducting a side-by-side comparison of the actual time elapsed and the corresponding time labels for both the treated and untreated ODE outputs, we gain valuable insights into the temporal shifts induced by the treatment. This comparative analysis enables us to quantify the acceleration of the healing process, providing a clear metric to assess the efficacy of the intervention.

For example, if the untreated model indicates that the peak macrophage activity (M1) is expected at day 3, while the treated model reaches this peak at day 2, this suggests a 33 precent reduction in time to reach this critical healing phase due to the treatment. Such data is invaluable, as it not only confirms the biological impact of the treatment but also aids in fine-tuning the dosages and timing of future therapeutic applications.

By integrating these findings with machine learning models, we can further refine the predictive accuracy of our treatment timelines. The non-linear RBF SVM model can adjust its parameters based on the treatment-induced changes, effectively recalibrating the day-based labels to match the new accelerated healing rates. This creates a robust, adaptable framework for real-time monitoring and management of wound healing, tailored to the individual responses to treatment observed in clinical practice.

5.6.2 Learning based on the ordinary differential equation data

The ODE model meticulously simulates the wound healing process, providing output probabilities for the key physiological stages: Hemostasis, Inflammatory, Proliferative, and Maturation. To translate these probabilities into actionable insights, we employ machine learning techniques to construct a predictive model. Specifically, a linear regression algorithm is designed to estimate the temporal progression of

healing by correlating stage probabilities with time.

This regression model is calibrated using data representative of average wound healing timelines. When fed data indicative of a more rapid healing process, the model interprets the advanced stage probabilities to suggest that the corresponding 'day' in the healing process is reached sooner than the average. In other words, the model's output will reflect a time label that progresses more quickly than real time, for wounds healing at an accelerated rate.

These predictions enable a direct comparison with the actual data derived from the ODE model, facilitating a nuanced analysis of the healing trajectory. Such comparisons are not merely academic; they hold practical significance in clinical decision-making. By quantifying the deviation from average healing times, clinicians can make informed decisions about the necessity for intervention, the adjustment of treatment plans, or the continuation of current therapeutic strategies.

Further, the integration of this regression model with real-time wound healing data allows for continuous refinement of the predictive algorithm. As new data is introduced, the model can adapt, adjusting its parameters to better reflect the diverse healing rates observed across different patient populations. This iterative process enhances the model's accuracy and reliability, leading to a robust tool that can be used to monitor wound healing progress, predict future stages, and guide clinical interventions tailored to individual patient needs.

In our approach, the Ordinary Differential Equations (ODE) model plays a pivotal role in quantifying the effect of treatments on the wound healing process. By simulating scenarios with and without treatment, we can discern the acceleration or deceleration caused by therapeutic interventions.

Consider a hypothetical situation to illustrate this: in the untreated scenario,

the model predicts that the wound is at day 3 of the healing process. When we introduce an 'actuator'—a proxy for the treatment—into the model, it simulates the wound's progress as if it were at day 4, indicating that the treatment has expedited the healing process.

To translate this model's insight into practical clinical assessments, we incorporate real-time data into a linear regression model. This model, trained on the probabilities output by the ODE model, is designed to predict the 'day' of healing a wound is experiencing. When we input real-world data into this regression model, the output provides us with a day-based estimate of the healing stage.

If the regression model predicts a healing day less than 3, it suggests that the wound is not as advanced in the healing process as the untreated model would predict—implying that the treatment may be hindering rather than helping recovery. Conversely, a prediction exceeding 3 days indicates that the treatment is beneficial, accelerating the healing beyond the baseline untreated scenario. An output matching 4 days would suggest that the wound is healing at an optimal rate, as projected by the treated model, following what we've defined as the 'ideal scenario'.

This method gives us a powerful decision-making tool: by comparing the regression model's predictions against the healing stages anticipated by the ODE model, clinicians can gauge the relative success of treatments. Such a comparison not only provides a snapshot of current healing status but also offers a predictive glimpse into the future of the wound's healing trajectory, allowing for adjustments in treatment strategy to align with the desired outcome. Ultimately, this enhances the potential for personalized wound care, tailored to the unique responses of each patient. Figure 5.4 illustrates the healing progression as predicted by ODE models. The graph plots ODE model predictions: simulating the natural healing process

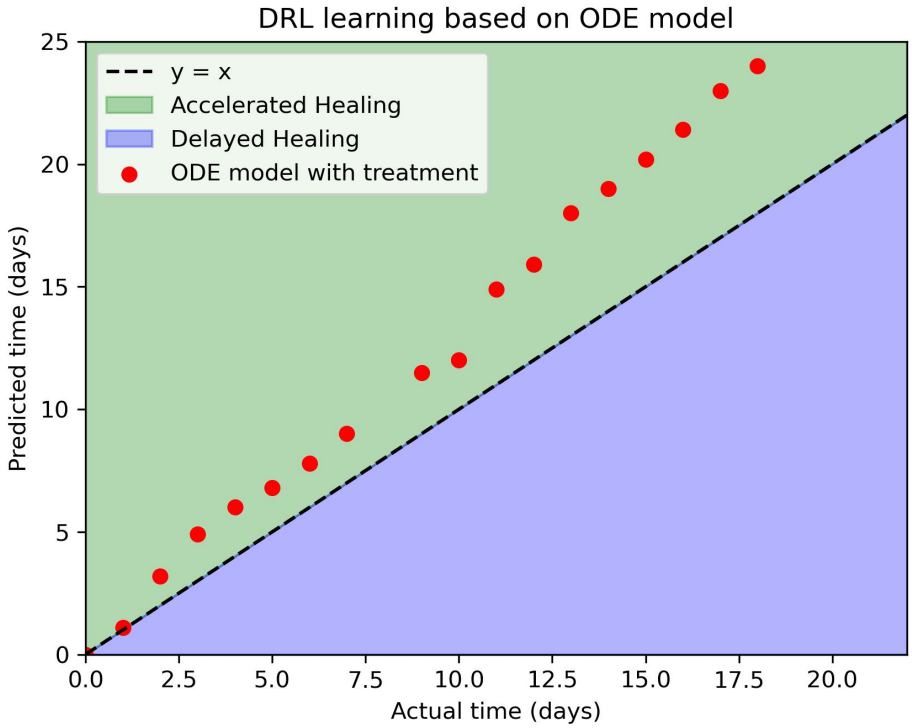

Figure 5.4: Deep Learning Training Reward based on ODE model, ODE with treatment. The accelerated healing area means positive reward, and the Delayed Healing area means negative reward.

without treatment (x=y line which is untreated), and another incorporating the accelerated healing due to treatment (treated model).

The ODE model incorporates treatment effects, predicting outcomes like accelerated or delayed healing based on various treatment actions. The DRL system is trained using a reward mechanism aligned with these outcomes. When the model predicts accelerated healing, the DRL agent receives a positive reward. This scenario implies that the chosen treatment strategy is effective, enhancing the healing process. Conversely, if the model predicts delayed healing, the DRL agent incurs a negative reward. This indicates a less effective treatment approach, hindering the healing

process.

The DRL agent's objective is to maximize the cumulative reward over time. Through iterative learning - a process of continuously adjusting its decisions and observing the resulting rewards - the agent learns the most effective strategies for applying treatment. It effectively understands:

Timing of intervention can be critical. The DRL agent learns the optimal times to administer treatment to maximize healing.

This includes understanding the appropriate dosages, frequencies, and types of treatment that lead to the best healing outcomes as per the ODE model's predictions.

By focusing on increasing the overall reward, the DRL agent is essentially learning to make decisions that align with the most favorable healing outcomes as predicted by the ODE model. This approach enables the DRL system to uncover strategies that might not be immediately obvious and could potentially offer new insights into effective treatment methodologies.

5.7 Deep Reinforcement Learning Model

In this model, I am using is a Deep Q-Network (DQN), which is a type of Deep Reinforcement Learning (DRL) algorithm. To understand this model and convergence we need to look at the fundamental components of the DQN and the mathematical principles behind them.

5.7.1 DQN

Deep Q-Networks (DQN) integrate deep learning with Q-learning, a type of reinforcement learning algorithm, to create robust decision-making agents. This document outlines the mathematical principles underlying DQN and discusses the conditions under which these models are expected to converge.

DQN comprises several key components including a neural network for function approximation, Q-learning for decision-making, exploration-exploitation balancing for effective learning, and experience replay for stable training.

In the implementation of the DQN, we used Python [101] as the programming language with Keras [29] and OpenAI Gym [102]. Keras provides a high-level neural networks API that is capable of running on top of TensorFlow, which is used here for building and training the neural network model. The DQN model is constructed using a Sequential model from Keras-RL [29], with Dense layers for neural network construction and ReLU and Linear activations for the hidden and output layers, respectively. The Adam optimizer is employed for training the network.

OpenAI Gym provides the environment for training and testing the agent. The rl package from Keras-RL [29] is utilized to define the DQN agent, with a Sequential Memory from Keras-RL [29] for experience replay and an Epsilon-Greedy policy for action selection. This setup is ideal for reinforcement learning tasks that require a balance between exploration and exploitation.

The model's training involves 5000 steps, with a warm-up phase of 10 steps and a target model update frequency of 0.01. This specific configuration is designed to efficiently learn from the environment 'wound-env' that is developed based on mathematical wound model. The details has been shown in the table 5.2.

Parameter	Value/Recommendation
Learning Rate	0.001
Optimizer	Adam
Batch Size	32
Minimum Epsilon	1 (decays to 0.001)
Epsilon Decay Factor	0.999
Steps before Training	10
Target Model Update Frequency	1000
Number of Hidden Layer Neurons	16
Activation Function for Hidden Layer	ReLU
Activation Function for Output Layer	Linear
Memory Limit	50,000
Number of Training Steps	5,000

Table 5.2: DQN Parameters for Highest Reward

5.7.2 Q-Learning Update Rule

Q-learning is an off-policy reinforcement learning algorithm used in DQN. The update rule for Q-learning is given by:

$$Q^{new}(s,a) = Q(s,a) + \alpha \left(r + \gamma \max_{a'} Q(s',a') - Q(s,a) \right) \tag{5.6}$$

where α is the learning rate, r is the reward, γ is the discount factor, s' is the new state, and a' are possible actions in the new state.

5.7.2.1 Loss Function for Neural Network

In DQN, a neural network approximates the Q-function. The loss function used to train this network is typically the mean squared error between the predicted Q-values and the target Q-values:

$$L(\theta) = \mathbb{E}\left[\left(r + \gamma \max_{a'} Q(s', a'; \theta^-) - Q(s, a; \theta)\right)^2\right] \tag{5.7}$$

Here, θ^- are the parameters of the target network, which are periodically updated with the parameters from the primary Q-network.

5.8 Reward Function

In our reinforcement learning model for the wound healing process, the reward function plays a critical role in guiding the agent's behavior. It is formulated to encourage actions that align with the optimal healing trajectory, as defined by the ODE model.

The reward function is structured to address three distinct scenarios in the healing process:

$$\text{Reward}(s, s') = \begin{cases} -1 & \text{if } s' < \text{normal healing stage} \\ \frac{s' - \text{normal healing stage}}{200} & \text{if } s' < 200 \\ 1 & \text{if } s' = 200 \end{cases} \tag{5.8}$$

In this formulation:

- s represents the current state.
- s' represents the next state.

- *normal healing stage* refers to a predefined benchmark in the healing process.

The reward is determined based on the progression of the healing process:

- A penalty of -1 is applied if s' is below the normal healing stage, indicating a lag in the healing process.
- A reward is computed as $\frac{s' - \text{normal healing stage}}{200}$ if s' is less than 200, reflecting ongoing healing.
- A reward of 1 is given when s' reaches day 20 which has been defined stage 200, marking the completion of the healing process.

5.8.1 Convergence of DQN

The convergence of my Deep Q-Network (DQN) implementation is influenced by several key factors, which are detailed as follows:

1. **Experience Replay**: The use of Sequential Memory from Keras-RL [29] for experience replay in DQN is a crucial factor for convergence. By storing past experiences and sampling random mini-batches from this memory for training, the DQN model breaks the correlation between consecutive learning updates. This approach leads to more stable and efficient learning as it prevents the model from overfitting to recent experiences and ensures a more comprehensive learning from diverse situations encountered in the environment.
2. **Stability through Fixed Q-Targets**: The DQN implementation typically involves using a separate target network for calculating the Q-value updates. This mechanism stabilizes the learning updates. If the same network was used for both predicting and updating Q-values, it could lead to oscillations or

divergence in the learning process. By using a slowly updated target network, my model maintains stability in learning, which is vital for convergence.

3. **Function Approximation with Neural Networks**: The neural network in DQN approximates the Q-function, which maps states and actions to expected rewards. The capability of neural networks to generalize from seen states to unseen states is crucial, especially in complex environments. My model successfully learned a generalized representation of the environment, this would significantly contribute to its ability to converge to an optimal policy.

4. **Empirical Indicators**: my DQN model demonstrated a decrease in loss over training epochs and an improvement in the agent's performance for example higher cumulative rewards, successful completion of healing in the wound simulation environment, these are strong empirical indicators of convergence. Such evidence suggests that the model was effectively learning the value function and improving its policy over time.

5. **Environment Complexity and Model Design**: The convergence also heavily depends on how well the DQN model's architecture and hyperparameters (like learning rate, discount factor, network structure) are suited to the specifics of the wound simulation environment. These were well-tuned to the complexity and the particular challenges of the environment, which facilitated effective learning and convergence.

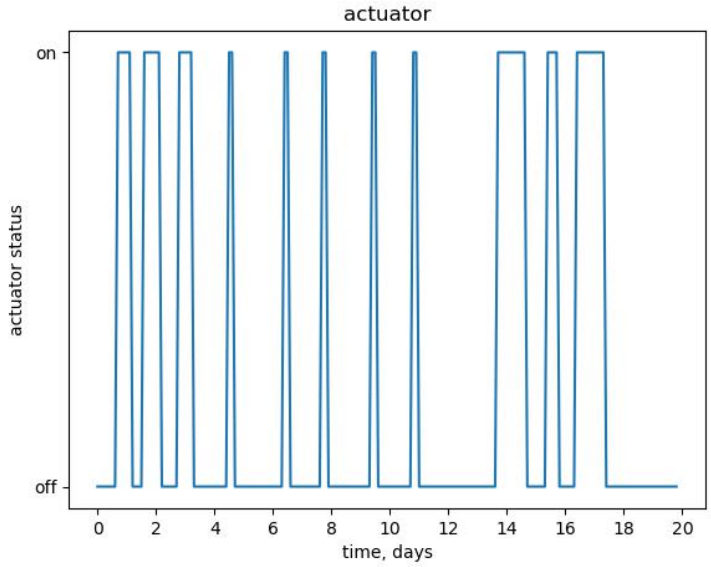

Figure 5.5: Drug treatment timing

5.9 Optimized Treatment

In this section, we delve into the application of a DRL model to optimize drug treatment regimens. We focus on the operational dynamics of an actuator designed to execute the drug administration policy derived through DRL algorithms. The actuator's role is to initiate or cease drug delivery in response to a policy that dictates the optimal timing and dosage based on a feedback from the environment. The actuator, as depicted in the 5.5, operates under a binary schema: an 'on' state where the drug is actively administered and an 'off' state indicating a pause in treatment. This binary operation is reflective of the discrete action space typically employed in DRL models, where the agent selects from a finite set of actions—in this case, to administer or withhold medication.

The functionality of the actuator is governed by a policy learned through a DRL framework. The DRL agent interacts with a simulated environment that models the patient's physiological response to drug treatment. Through continuous

interaction and feedback, characterized by a reward system, the DRL agent learns an optimal policy that maximizes therapeutic outcomes while minimizing adverse effects.

The optimization process involves the agent evaluating the consequences of actions taken at various states, informed by both immediate and future estimated rewards. As the agent traverses the state-action space, it updates its policy to favor actions that lead to higher cumulative rewards, which align with treatment objectives such as efficacy, safety, and adherence to clinical guidelines.

The pattern observed in the figure suggests that the DRL model has identified a treatment schedule that alternates between drug administration and withholding. This intermittent dosing strategy could be aimed at preventing drug resistance, managing side effects, or aligning with the drug's pharmacokinetic properties. The convergence of the DRL model to this pattern indicates successful learning, which is validated by empirical evidence such as improved patient outcomes or adherence to optimal treatment pathways.

The actuator's operation reflects a nuanced understanding of the treatment landscape, indicating that DRL can effectively handle complex decision-making processes. The convergence of the DRL model to an optimal policy, as evidenced by the figure, underscores the model's ability to navigate the intricate balance between various treatment factors.

Chapter 6

Sliding Mode Controller

6.1 Introduction

A sliding mode controller is a type of nonlinear control system used to regulate the behavior of a dynamic system [103] It is designed to work in the presence of model uncertainties, external disturbances, and other types of nonlinearities. The basic idea behind a sliding mode controller is to create a sliding surface, which is a hyperplane in the state space of the system, that the system will "slide" along to reach a desired state. This sliding surface is designed such that when the system is away from the desired state, the sliding mode controller generates a control signal that drives the system towards the sliding surface. Once the system reaches the sliding surface, the controller generates a different control signal that keeps the system on the sliding surface, and hence at the desired state. The sliding mode controller uses a discontinuous control signal to force the system onto the sliding surface. This means the control signal switches between two or more values depending on the system's position in the state space. The discontinuous control signal is designed such that it does not excite high-frequency oscillations in the system. Sliding mode control has

been applied to a wide variety of systems, including electrical circuits, mechanical systems, and aerospace systems. Its robustness to uncertainties and disturbances makes it a popular choice in many applications [104]. Sliding mode control has also been applied in various biological systems, including the control of gene expression, the regulation of metabolic networks, and the control of glucose levels in diabetic patients [105]. Sliding mode control has shown promise in the control of various biological systems, providing a robust and effective control strategy in the presence of uncertainties and disturbances.

6.2 Sliding mode controller for wound healing

Sliding mode control (SMC) is a form of control theory that is particularly suited to managing systems characterized by nonlinearities and uncertainties, which makes it an excellent candidate for regulating the complex biological processes involved in wound healing. This method treats the healing wound as a dynamic system, where various cellular interactions and the effects of growth factors are the driving forces of recovery. In the application of SMC to wound healing, the 'sliding surface'—a mathematical construct representing the desired system behavior—is carefully designed to correspond with the stages of wound recovery, with the ultimate objective of minimizing the time required for a wound to heal.

In the realm of potential therapies, electrical stimulation has garnered attention due to its ability to accelerate wound healing. The mechanism by which it operates involves the application of low-level electrical currents to the wound site, which can be administered directly or through electrodes positioned nearby. This technique has been found to foster various healing-enhancing biological responses, such

as promoting cell migration, proliferation, and differentiation, alongside increasing blood flow and angiogenesis.

Another therapeutic avenue being explored is the use of pharmacological agents like fluoxetine, which is known to possess antioxidant and anti-inflammatory properties. These characteristics are believed to help reduce the oxidative stress and inflammation that can impede wound healing. Additionally, fluoxetine has been associated with the promotion of angiogenesis and collagen synthesis, both of which play a critical role in the repair and closure of wounds.

The advantage of SMC lies in its robust performance amidst the variability inherent in biological systems like wound healing. Traditional control strategies may fall short when faced with such complexity, but the adaptability of SMC allows for precise management even when faced with unpredictable changes in the system's dynamics.

For instance, in the case of electrical stimulation therapy, SMC can adeptly manage variations in wound resistance and other factors that may influence the voltage requirements. By focusing on maintaining a consistent current, the SMC ensures that the wound receives the necessary level of stimulation conducive to healing. This is accomplished by adjusting the voltage in real-time to maintain the target current, thereby ensuring the therapeutic benefits are delivered reliably throughout the treatment duration.

When considering drug delivery, SMC can be instrumental in maintaining the desired drug concentration at the wound site. It achieves this by regulating the drug release in response to the monitored current, ensuring that the actual concentration aligns with the therapeutic requirements. If the wound's electrical properties change—possibly indicating changes in tissue absorption—the SMC is

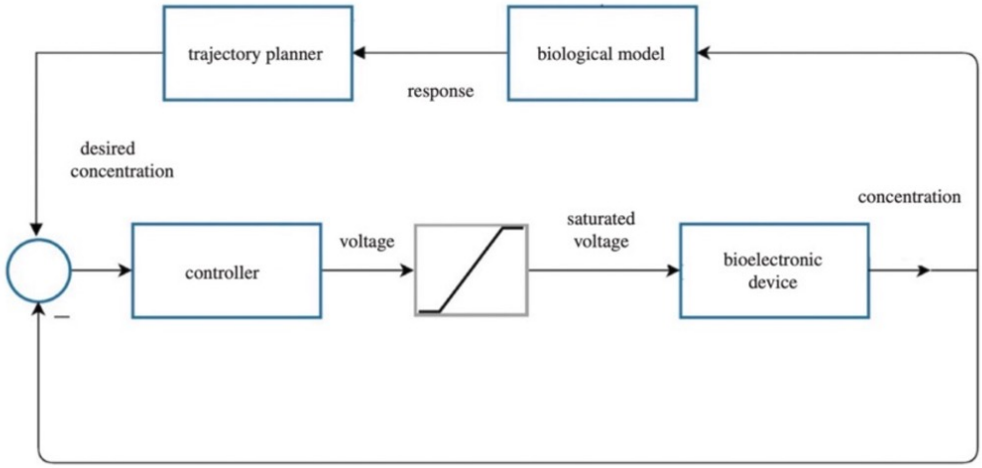

Figure 6.1: Scheme of sliding mode controller.

capable of modifying the drug delivery rate to preserve the required drug concentration at the site of the wound.

Figure 6.1 depicts the complex architecture of the Sliding Mode Control (SMC) system, engineered specifically for this study on wound healing. The diagram illustrates the synergy between various components of the SMC system, including sensors to monitor wound parameters, the control algorithm that governs the system's response, actuators responsible for administering treatment, and the feedback mechanisms that ensure treatment efficacy. This integrated system is designed to administer precise dosages of therapeutic agents or electrical currents, modulating the wound environment to optimize the healing process. The sophistication of the SMC system lies in its ability to maintain the desired therapeutic conditions despite the inherent uncertainties and variabilities of biological systems.

6.2.1 Sliding mode controller for the device with 8 channels

In the landscape of modern wound management, the intricacies of the healing environment demand a nuanced approach to treatment. This is where multi-channel SMC comes into play, elevating the standard of care by offering bespoke treatment modalities. A typical advanced wound care device may feature eight distinct channels, each of which can be independently controlled by a dedicated SMC system. The rationale behind such a design is to address the unique conditions of various wound segments that may differ in terms of size, depth, severity, and healing stage.

Each channel's independent SMC allows for the precise application of electrical currents tailored to the specific requirements of different wound areas. For instance, parts of a wound that are deeper or more severely affected may necessitate a different current intensity compared to more superficial or less severe areas. The SMC's ability to independently modulate these variables ensures that each section of the wound receives optimal treatment without the risk of overstimulation or under-treatment.

The independent channels also enhance the adaptability of the device. As the wound heals and its characteristics evolve, each SMC can adjust the treatment parameters in real-time, ensuring that the therapeutic action remains aligned with the current healing stage. This dynamic adjustment is crucial, especially in wounds with irregular shapes or those exhibiting heterogeneous healing patterns.

Moreover, the segmented control mechanism inherently enhances the safety and reliability of the device. In the event of a channel malfunction or anomaly, the system is designed to isolate the issue, allowing the unaffected channels to continue

their function seamlessly. This ensures that treatment is not interrupted, maintaining the continuity and effectiveness of the healing process.

6.3 Chicken Breast Sliding Mode Controller (SMC) test

In the development of cutting-edge wound care technology, empirical testing plays an indispensable role in validating the functionality and dependability of the device. For our SMC-equipped wound care device, the preliminary testing was meticulously carried out using a chicken breast as a proxy for human tissue. This initial experiment spanned three days and was instrumental in calibrating the device's performance under varied conditions reflective of real-world scenarios.

The chicken breast model was chosen due to its tissue consistency and similarity to human flesh, providing a relevant medium for testing. During this crucial phase, the device underwent rigorous trials to ensure the SMC controller consistently maintained the target treatment values, such as drug concentration or electrical current, across different test states. This step was vital to detect and rectify any discrepancies in device behavior before proceeding to more complex models.

Throughout the testing period, the research team encountered and overcame several challenges. Notably, interruptions in the wifi connectivity highlighted a vulnerability in the system's ability to maintain a closed-loop control. In response, the team innovated the communication protocols between the controller and the printed circuit board (PCB) code. By streamlining the data transfer and integrating the closed-loop controller with the device controller code, the team effectively reduced the computational burden and enhanced the system's stability.

These modifications significantly improved the device's resilience, ensuring

that it could reliably administer the intended treatment even in the event of connectivity issues. The iterative process of testing, identifying faults, and refining the system was pivotal in advancing the device from a prototype to a reliable instrument capable of delivering consistent care.

As a consequence of these meticulous efforts, the device's performance was optimized, paving the way for subsequent testing phases involving more complex biological models such as pigs. This step-by-step escalation in testing complexity is essential, laying a solid foundation for the eventual goal of clinical application. The insights gained from the chicken breast tests were crucial, not only for immediate improvements but also for informing the future direction of the research, ultimately contributing to the development of more effective wound healing modalities.

6.4 Testing the SMC controller for wound care on Pig

Sliding Mode Control (SMC) technology is at the forefront of enhancing wound care by administering meticulously controlled electrical currents to promote faster healing and improved patient recovery. Following promising initial tests on chicken breast tissue, which confirmed the SMC system's functionality and allowed for the rectification of technical issues, the research advanced to a more complex biological model—wound care on a pig.

Testing the SMC system on porcine wounds was a decisive step, bridging the gap between controlled laboratory conditions and the variable circumstances of live biological systems. The pig's wound offered a realistic platform for the SMC system to adapt its operations to the nuanced changes in a live wound environment, showcasing its ability to tailor the current to the wound's specific healing requirements.

During the practical trials, the SMC demonstrated its potential by accurately managing the current over brief intervals, aligning with predefined treatment durations. This experiment yielded critical data, informing subsequent iterations of the device design to enhance its therapeutic efficacy.

Despite the system's successes, challenges arose when attempting to sustain high targeted currents over extended periods. Figure 3 illustrates these findings through an error plot, revealing the variance between the target current and the actual current delivered by the device. The graph indicates a consistent adherence to the target current for approximately 3.5 hours before the actual current levels began to wane. This discrepancy was attributed to hardware constraints and unforeseen technical difficulties, which are being actively addressed by the hardware development team.

The testing phase on the pig wound was not just a technical evaluation, but a valuable learning experience, revealing both the capabilities and current limitations of the SMC system. These insights are vital for driving hardware enhancements and ensuring that future implementations of the SMC system can maintain consistent treatment levels throughout the healing process. Such advancements are essential for moving closer to the goal of a fully reliable and effective wound care treatment system, one that can adapt in real-time to the changing demands of wound healing and ultimately improve patient outcomes. In the practical application of the Sliding Mode Control (SMC) for wound care, the controller faced challenges in consistently maintaining the high targeted current necessary for optimal wound treatment. This was evident in the data presented in Figure 6.2 , which provides an current plot showing the targeted and the actual currents delivered by the device.

The slight variations in current delivery observed in this study suggest that

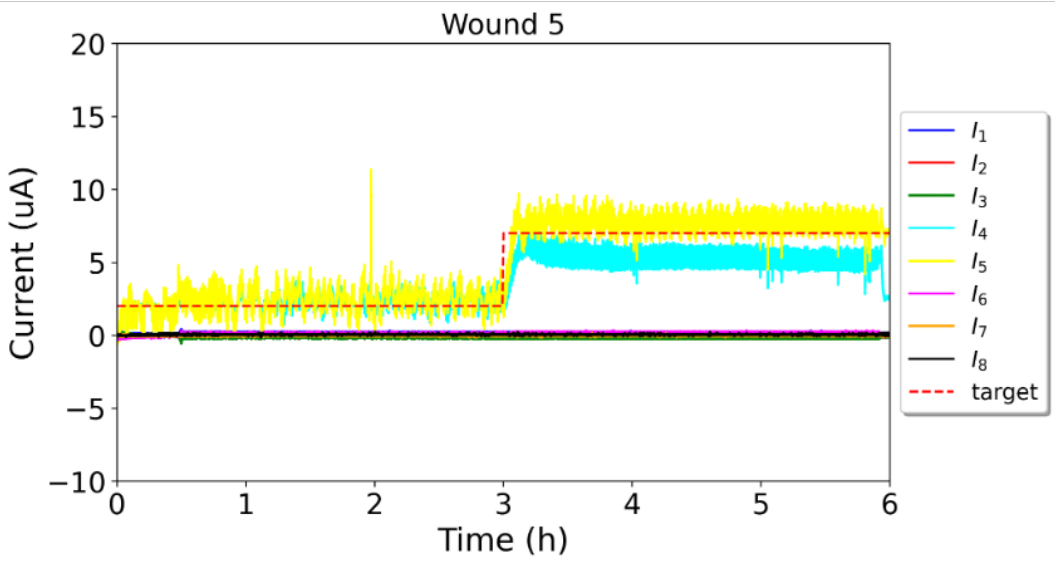

Figure 6.2: Tracking Device and Target Currents Over Time: Target Shifts from 3 µA to 7 µA at 3 Hours

while the system is generally effective, there is room for improvement to enhance the robustness of electrical therapies for wound healing. These minor deviations may be due to subtle hardware constraints, such as the power supply's output or the precision of the current delivery mechanism. Incremental hardware enhancements could lead to a more consistently accurate adherence to the target current, ensuring a more reliable and stable treatment delivery.

The hardware team is dedicated to addressing these challenges by delving into the root causes of the observed discrepancies. Potential solutions may involve redesigning certain hardware components to handle higher loads, improving the stability and responsiveness of the control algorithm, or enhancing the overall system design to prevent such changes in current.

As these improvements are conceptualized and implemented, a continuous feedback loop from testing will inform the iterative development process. With each round of testing and refinement, the device moves closer to achieving a level of

reliability and precision that meets the stringent requirements of medical treatment devices. The ultimate objective remains clear: to develop a wound care system that delivers consistent, reliable, and effective treatment to enhance the healing process and improve patient outcomes.

The journey from prototype to a reliable medical device is often a complex one, marked by challenges and learnings. The case of the SMC system for wound care is no different. Through diligent testing, problem-solving, and engineering, the team strives to perfect a device that can revolutionize wound management and set a new standard in therapeutic care.

The experimental phase utilizing the SMC system for wound healing is more than a test of current capabilities—it's an invaluable learning experience that reveals the system's current limitations and potential enhancements. The insights garnered from this phase are instrumental for the engineering team, who are now equipped with real-world data to refine and optimize the device's performance, ensuring that it can stand up to the unpredictable and varied conditions of clinical use.

During the testing phase, the SMC system was subjected to a range of scenarios, each designed to challenge the controller's adaptability and responsiveness. This rigorous evaluation not only confirmed the system's potential but also highlighted areas that required redesign to better serve the needs of patients, particularly those with complex wound profiles that demand a higher degree of care.

The process of refining the SMC system is iterative and patient-centric. With each iteration, the device inches closer to the ideal of providing tailored, precise treatment. Adjustments to the hardware and software are made with the ultimate goal of delivering consistent and effective therapy, addressing the unique healing dynamics of each wound.

Moreover, this phase of development underscores the importance of adaptability in medical device technology. As the SMC system is fine-tuned, it becomes more capable of accommodating a broader spectrum of wound types, severities, and patient-specific factors. It's a journey towards creating a device that not only treats wounds but does so with an unprecedented level of precision and personalization.

6.5 Conclusion

In conclusion, the testing and subsequent refinement of the SMC controller are integral to its evolution. This process is critical for ensuring that the final product is not only effective in controlled trials but also successful in the hands of medical professionals treating real wounds. By methodically addressing each challenge and incorporating the feedback into the device's design, the team is committed to delivering a wound care solution that enhances healing outcomes, reduces recovery times, and ultimately improves the quality of life for patients.

Bibliography

[1] Sabine A Eming, Paul Martin, and Marjana Tomic-Canic. Wound repair and regeneration: Mechanisms, signaling, and translation. *Science translational medicine*, 6(265):265sr66–265sr66, 2014.

[2] Peng Zhao, Wei He, Yanyan Liu, and Gang Zhao. The role of neurites in wound healing: A review. *Frontiers in Bioengineering and Biotechnology*, 7:189, 2019.

[3] Prakash Monika, Mathikere Naganna Chandraprabha, Annapoorni Rangarajan, P. Veena Waiker, and Kotamballi N. Chidambara Murthy. Challenges in healing wound: Role of complementary and alternative medicine. *Frontiers in Nutrition*, 8, 2022.

[4] R. S. Kirsner, L. Bello, and O. Kerdel. Chronic wounds: Evaluation and management. *American Family Physician*, 99(4):219–226, 2019.

[5] R. Levi-Montalcini. The nerve growth factor 35 years later. *Science*, 237(4819):1154–1162, 1987.

[6] Z. Liu, H. Wu, and S. Huang. Role of ngf and its receptors in wound healing (review). *Experimental and Therapeutic Medicine*, 21:263, 2021.

[7] Mahmoud G. El Baassiri, Laura Dosh, Hanine Haidar, Alice Gerges, Silma Baassiri, Angelo Leone, Francesca Rappa, and Abdo Jurjus. Nerve growth

factor and burn wound healing: Update of molecular interactions with skin cells. *Burns*, 49:989–1002, 2023.

[8] Ji-Cai Chen, Bei-Bei Lin, Hou-Wen Hu, Cai Lin, Wen-Yang Jin, Fa-Biao Zhang, Yan-An Zhu, Cai-Jiao Lu, Xiao-Jie Wei, and Rui-Jie Chen. Ngf accelerates cutaneous wound healing by promoting the migration of dermal fibroblasts via the pi3k/akt-rac1-jnk and erk pathways. *BioMed Research International*, 2014:Article ID 547187, 2014.

[9] Mariya Barch, Melanie M. Cobb, Zachary Tokuno, Jennifer Leddy, Keili Prenton, Linus Manubens-Gil, Nicole Bellini, Stephanie Lam, Julia A. Kaye, Mara Dierssen, and Steven Finkbeiner. A deep learning approach to neurite prediction in high throughput fluorescence imaging. *bioRxiv*, 2021.

[10] Giorgio A. Ascoli, Duncan E. Donohue, and Maryam Halavi. Neuromorpho.org: A central resource for neuronal morphologies. *Journal of Neuroscience*, 27(35):9247–9251, 2007.

[11] M. Nishio, C. Nagashima, S. Hirabayashi, A. Ohnishi, K. Sasaki, T. Sagawa, M. Hamada, and T. Yamashita. Convolutional auto-encoder for image denoising of ultra-low-dose ct. *Heliyon*, 3(8), 2017.

[12] D. Queen. Artificial intelligence and machine learning in wound care-the wounded machine! *International Wound Journal*, 16(2):311, 2019.

[13] Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. Reinforcement learning in healthcare: A survey. *ACM Comput. Surv.*, 55(1), nov 2021.

[14] NX Landén, D Li, and M Ståhle. Transition from inflammation to proliferation:

a critical step during wound healing. *Cellular and Molecular Life Sciences*, 73(20):3861–3885, 2016.

[15] N. Noskovicova, B. Hinz, and P. Pakshir. Implant fibrosis and the underappreciated role of myofibroblasts in the foreign body reaction. *Cells*, 10(7):1794, 2021.

[16] P. Kolimi, S. Narala, D. Nyavanandi, AAA Youssef, and N. Dudhipala. Innovative treatment strategies to accelerate wound healing: Trajectory and recent advancements. *Cells*, 11(15):2439, 2022.

[17] MA Weigelt, HA Lev-Tov, M Tomic-Canic, WD Lee, R Williams, D Strasfeld, RS Kirsner, and IM Herman. Advanced wound diagnostics: Toward transforming wound care into precision medicine. *Advances in Wound Care*, 11(6):330–359, 2022.

[18] AS Gross, AC Harry, CS Clifton, and O Della Pasqua. Clinical trial diversity: An opportunity for improved insight into the determinants of variability in drug response. *British Journal of Clinical Pharmacology*, 88(6):2700–2717, 2022.

[19] M. Berezo, J. Budman, D. Deutscher, CT Hess, K. Smith, and D. Hayes. Predicting chronic wound healing time using machine learning. *Advances in Wound Care*, 11(6):281–296, 2022.

[20] L. Chen, Y. W. Zhang, and S. C. Zhang. Optimal drug dosage control strategy of immune systems using reinforcement learning. *IEEE Access*, 11:1269–1279, 2023.

[21] A.S. Mehta, S. Teymoori, C. Recendez, et al. Quantifying innervation facilitated by deep learning in wound healing. *Scientific Reports*, 13:16885, 2023.

[22] Mohammed Ashrafi, Mohamed Baguneid, and Ardeshir Bayat. The role of neuromediators and innervation in cutaneous wound healing. *Acta dermato-venereologica*, 96(5):587–594, 2016.

[23] Michelle D Bagood and R Rivkah Isseroff. Trpv1: Role in skin and skin diseases and potential target for improving wound healing. *International journal of molecular sciences*, 22(11):6135, 2021.

[24] Vincent Falanga, Roslyn Rivkah Isseroff, Athena M Soulika, Marco Romanelli, David Margolis, Suzanne Kapp, Mark Granick, and Keith Harding. Chronic wounds. *Nature Reviews Disease Primers*, 8(1):50, 2022.

[25] Lars H Evers, Dhaval Bhavsar, and Peter Mailänder. The biology of burn injury. *Experimental dermatology*, 19(9):777–783, 2010.

[26] Christian M Kähler, Manfred Herold, Norbert Reinisch, and Christian J Wiedermann. Interaction of substance p with epidermal growth factor and fibroblast growth factor in cyclooxygenase-dependent proliferation of human skin fibroblasts. *Journal of cellular physiology*, 166(3):601–608, 1996.

[27] Benedetta Brazzini, Ilaria Ghersetich, Jana Hercogova, and Torello Lotti. The neuro-immuno-cutaneous-endocrine network: relationship between mind and skin. *Dermatologic Therapy*, 16(2):123–131, 2003.

[28] Daniel R Fregoso, Yasmin Hadian, Anthony C Gallegos, Doniz Degovics, John Maaga, Ciara E Keogh, Isaiah Kletenik, Melanie G Gareau, and R Rivkah Isseroff. Skin-brain axis signaling mediates behavioral changes after skin wounding. *Brain, Behavior, & Immunity-Health*, 15:100279, 2021.

[29] Jérémy Chéret, Nicolas Lebonvallet, Jean-Luc Carré, Laurent Misery, and Christelle Le Gall-Ianotto. Role of neuropeptides, neurotrophins, and neurohormones in skin wound healing. *Wound repair and regeneration*, 21(6):772–788, 2013.

[30] Allison R Barker, Gedge D Rosson, and A Lee Dellon. Wound healing in denervated tissue. *Annals of plastic surgery*, 57(3):339–342, 2006.

[31] Michael Schäffer, Thomas Beiter, Horst Dieter Becker, and Thomas K Hunt. Neuropeptides: mediators of inflammation and tissue repair? *Archives of surgery*, 133(10):1107–1116, 1998.

[32] Toshihiro Tanaka, Kiichiro Danno, Kouichi Ikai, and Sadao Imamura. Effects of substance p and substance k on the growth of cultured keratinocytes. *Journal of investigative dermatology*, 90(3), 1988.

[33] Tai-Ping D Fan, De-En Hu, Steven Guard, G Austin Gresham, and Keith J Watling. Stimulation of angiogenesis by substance p and interleukin-1 in the rat and its inhibition by nk1 or interleukin-1 receptor antagonists. *British journal of pharmacology*, 110(1):43–49, 1993.

[34] Pornprom Muangman, Richard N Tamura, Lara A Muffley, F Frank Isik, Jeffrey R Scott, Chengyu Xie, Gary Kegel, Stephen R Sullivan, Zhi Liang, and Nicole S Gibran. Substance p enhances wound closure in nitric oxide synthase knockout mice. *Journal of Surgical Research*, 153(2):201–209, 2009.

[35] Marina Ziche, Lucia Morbidelli, Marco Pacini, Piero Dolara, and Carlo Alberto Maggi. Nk1-receptors mediate the proliferative response of human fibroblasts to tachykinins. *British journal of pharmacology*, 100(1):11–14, 1990.

[36] V Altun, TE Hakvoort, PPM Van Zuijlen, Th H Van Der Kwast, and EP Prens. Nerve outgrowth and neuropeptide expression during the remodeling of human burn wound scars: A 7-month follow-up study of 22 patients. *Burns*, 27(7):717–722, 2001.

[37] Cory A Dunnick, Nicole S Gibran, and David M Heimbach. Substance p has a role in neurogenic mediation of human burn wound healing. *The Journal of burn care & rehabilitation*, 17(5):390–396, 1996.

[38] MD Basson and RE Burney. Defective wound healing in patients with paraplegia and quadriplegia. *Surgery, gynecology & obstetrics*, 155(1):9–12, 1982.

[39] J Marin, J Nixon, and C Gorecki. A systematic review of risk factors for the development and recurrence of pressure ulcers in people with spinal cord injuries. *Spinal cord*, 51(7):522–527, 2013.

[40] Mark H Longair, Dean A Baker, and J Douglas Armstrong. Simple neurite tracer: open source software for reconstruction, visualization and analysis of neuronal processes. *Bioinformatics*, 27(17):2453–2454, 2011.

[41] Curtis T Rueden, Johannes Schindelin, Mark C Hiner, Barry E DeZonia, Alison E Walter, Ellen T Arena, and Kevin W Eliceiri. Imagej2: Imagej for the next generation of scientific image data. *BMC bioinformatics*, 18:1–26, 2017.

[42] Nan Jiang, Jeffrey P Rasmussen, Joshua A Clanton, Marci F Rosenberg, Kory P Luedke, Mark R Cronan, Edward D Parker, Hyeon-Jin Kim, Joshua C Vaughan, Alvaro Sagasti, et al. A conserved morphogenetic mechanism for epidermal ensheathment of nociceptive sensory neurites. *Elife*, 8:e42455, 2019.

[43] John C Gensel, David L Schonberg, Jessica K Alexander, Dana M McTigue, and Phillip G Popovich. Semi-automated sholl analysis for quantifying changes in growth and differentiation of neurons and glia. *Journal of neuroscience methods*, 190(1):71–79, 2010.

[44] Robert A Underwood, Nicole S Gibran, Lara A Muffley, Marcia L Usui, and John E Olerud. Color subtractive–computer-assisted image analysis for quantification of cutaneous nerves in a diabetic mouse model. *Journal of Histochemistry & Cytochemistry*, 49(10):1285–1291, 2001.

[45] Gigi J Ebenezer, Peter Hauer, Christopher Gibbons, Justin C McArthur, and Michael Polydefkis. Assessment of epidermal nerve fibers: a new diagnostic and predictive tool for peripheral neuropathies. *Journal of Neuropathology & Experimental Neurology*, 66(12):1059–1073, 2007.

[46] Dimitar Evdokimov, Philine Dinkel, Johanna Frank, Claudia Sommer, and Nurcan Üçeyler. Characterization of dermal skin innervation in fibromyalgia syndrome. *PLoS One*, 15(1):e0227674, 2020.

[47] Peter YK Van den Bergh, Robert DM Hadden, Pierre Bouche, David R Cornblath, Angelika Hahn, Isabel Illa, Carol L Koski, J-M Léger, Eduardo Nobile-Orazio, John Pollard, et al. European federation of neurological societies/peripheral nerve society guideline on management of chronic inflammatory demyelinating polyradiculoneuropathy: report of a joint task force of the european federation of neurological societies and the peripheral nerve society—first revision. *European journal of neurology*, 17(3):356–363, 2010.

[48] Giuseppe Lauria, Mayienne Bakkers, Christoph Schmitz, Raffaella Lombardi,

Paola Penza, Grazia Devigili, A Gordon Smith, Sung-Tsieh Hsieh, Svein I Mellgren, Thirugnanam Umapathi, et al. Intraepidermal nerve fiber density at the distal leg: a worldwide normative reference study. *Journal of the Peripheral Nervous System*, 15(3):202–207, 2010.

[49] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017.

[50] Kyungsang Kim, Dufan Wu, Kuang Gong, Joyita Dutta, Jong Hoon Kim, Young Don Son, Hang Keun Kim, Georges El Fakhri, and Quanzheng Li. Penalized pet reconstruction using deep learning prior and local linear fitting. *IEEE transactions on medical imaging*, 37(6):1478–1487, 2018.

[51] Hsin-ya Yang, Fernando Fierro, Michelle So, Daniel J Yoon, Alan Vu Nguyen, Anthony Gallegos, Michelle D Bagood, Tomas Rojo-Castro, Alan Alex, Heather Stewart, et al. Combination product of dermal matrix, human mesenchymal stem cells, and timolol promotes diabetic wound healing in mice. *Stem cells translational medicine*, 9(11):1353–1364, 2020.

[52] Robert D Galiano, Joseph Michaels, V, Michael Dobryansky, Jamie P Levine, and Geoffrey C Gurtner. Quantitative and reproducible murine model of excisional wound healing. *Wound repair and regeneration*, 12(4):485–492, 2004.

[53] Shin Ae Park, Leandro BC Teixeira, Vijay Krishna Raghunathan, Jill Covert, Richard R Dubielzig, Roslyn Rivkah Isseroff, Michael Schurr, Nicholas L Abbott, Jonathan McAnulty, and Christopher J Murphy. Full-thickness splinted skin

wound healing models in db/db and heterozygous mice: implications for wound healing impairment. *Wound Repair and Regeneration*, 22(3):368–380, 2014.

[54] BG McCarthy, S-T Hsieh, A Stocks, P Hauer, C Macko, DR Cornblath, JW Griffin, and JC McArthur. Cutaneous innervation in sensory neuropathies: evaluation by skin biopsy. *Neurology*, 45(10):1848–1855, 1995.

[55] Haitao Shen, Pamela Yao, Eunyoung Lee, David Greenhalgh, and Athena M Soulika. Interferon-gamma inhibits healing post scald burn injury. *Wound repair and regeneration*, 20(4):580–591, 2012.

[56] Abijeet Singh Mehta, Prajakta Deshpande, Anuradha Venkatakrishnan Chimata, Panagiotis A Tsonis, and Amit Singh. Newt regeneration genes regulate wingless signaling to restore patterning in drosophila eye. *Iscience*, 24(10), 2021.

[57] Jean Jew, Evelyn Berger, Richard Berger, and Yu-Te Lin. Fluorescence immunohistochemistry and confocal scanning laser microscopy a protocol for studies of joint innervation. *Acta Orthopaedica Scandinavica*, 74(6):689–696, 2003.

[58] Anuradha Venkatakrishnan Chimata, Prajakta Deshpande, Abijeet Singh Mehta, and Amit Singh. Protocol to study cell death using tunel assay in drosophila imaginal discs. *STAR protocols*, 3(1):101140, 2022.

[59] Chuong Minh Nguyen, Danielle Marie Tartar, Michelle Dawn Bagood, Michelle So, Alan Vu Nguyen, Anthony Gallegos, Daniel Fregoso, Jorge Serrano, Duc Nguyen, Doniz Degovics, et al. Topical fluoxetine as a novel therapeutic that improves wound healing in diabetic mice. *Diabetes*, 68(7):1499–1507, 2019.

[60] Abijeet Singh Mehta, Pin Ha, Kan Zhu, ShiYu Li, Kang Ting, Chia Soo, Xinli Zhang, and Min Zhao. Physiological electric fields induce directional migration of mammalian cranial neural crest cells. *Developmental biology*, 471:97–105, 2021.

[61] Abijeet Singh Mehta, Brijesh K Singh, Nandita Singh, D Archana, Kirti Snigdha, Robert Harniman, Sameer S Rahatekar, RP Tewari, and PK Dutta. Chitosan silk-based three-dimensional scaffolds containing gentamicin-encapsulated calcium alginate beads for drug administration and blood compatibility. *Journal of biomaterials applications*, 29(9):1314–1325, 2015.

[62] Nathalie Van Acker, Michael Ragé, Ellen Sluydts, Michiel WM Knaapen, Martine De Bie, Maarten Timmers, Erik Fransen, Carla Duymelinck, Stefanie De Schepper, Praveen Anand, et al. Automated pgp9. 5 immunofluorescence staining: a valuable tool in the assessment of small fiber neuropathy? *BMC Research Notes*, 9:1–11, 2016.

[63] C-J Dalsgaard, M Rydh, and A Haegerstrand. Cutaneous innervation in man visualized with protein gene product 9.5 (pgp 9.5) antibodies. *Histochemistry*, 92(5):385–390, 1989.

[64] Lixin Wang, Marita Hilliges, Tomas Jernberg, Desirée Wiegleb-Edström, and Olle Johansson. Protein gene product 9.5-immunoreactive nerve fibres and cells in human skin. *Cell and tissue research*, 261:25–33, 1990.

[65] Vineeth Murali and PV Sudeep. Image denoising using dncnn: an exploration study. In *Advances in Communication Systems and Networks: Select Proceedings of ComNet 2019*, pages 847–859. Springer, 2020.

[66] Ademola E Ilesanmi and Taiwo O Ilesanmi. Methods for image denoising using convolutional neural network: a review. *Complex & Intelligent Systems*, 7(5):2179–2198, 2021.

[67] Shubhankar Rawat, KPS Rana, and Vineet Kumar. A novel complex-valued convolutional neural network for medical image denoising. *Biomedical Signal Processing and Control*, 69:102859, 2021.

[68] Chengdi Wang, Xiuyuan Xu, Jun Shao, Kai Zhou, Kefu Zhao, Yanqi He, Jingwei Li, Jixiang Guo, Zhang Yi, Weimin Li, et al. Deep learning to predict egfr mutation and pd-l1 expression status in non-small-cell lung cancer on computed tomography images. *Journal of Oncology*, 2021, 2021.

[69] Mathieu Blais, Lorène Mottier, Marie-Anne Germain, Sabrina Bellenfant, Sébastien Cadau, and François Berthod. Sensory neurons accelerate skin reepithelialization via substance p in an innervated tissue-engineered wound healing model. *Tissue Engineering Part A*, 20(15-16):2180–2188, 2014.

[70] Peter G Smith and Manxi Liu. Impaired cutaneous wound healing after sensory denervation in developing rats: effects on cell proliferation and apoptosis. *Cell and tissue research*, 307:281–291, 2002.

[71] JA Ellis, H Goldstein, and CJ Winfree. Nerve repair. 2014.

[72] Abijeet S Mehta and Amit Singh. Insights into regeneration tool box: An animal model approach. *Developmental biology*, 453(2):111–129, 2019.

[73] Abijeet Singh Mehta, Agustin Luz-Madrigal, Jian-Liang Li, Panagiotis A Tsonis, and Amit Singh. Comparative transcriptomic analysis and structure prediction of novel newt proteins. *PloS one*, 14(8):e0220416, 2019.

[74] Yang Jiao, Yu-Wan Liu, Wei-Gong Chen, and Jing Liu. Neuroregeneration and functional recovery after stroke: Advancing neural stem cell therapy toward clinical application. *Neural regeneration research*, 16(1):80–92, 2021.

[75] Tingting Weng, Pan Wu, Wei Zhang, Yurong Zheng, Qiong Li, Ronghua Jin, Haojiao Chen, Chuangang You, Songxue Guo, Chunmao Han, et al. Regeneration of skin appendages and nerves: current status and further challenges. *Journal of translational medicine*, 18(1):1–17, 2020.

[76] Betty Laverdet, Aurore Danigo, Dorothée Girard, Laurent Magy, Claire Demiot, and Alexis Desmoulière. Skin innervation: important roles during normal and pathological cutaneous repair. 2015.

[77] Anil Sebastian, Susan W Volk, Poonam Halai, James Colthurst, Ralf Paus, and Ardeshir Bayat. Enhanced neurogenic biomarker expression and reinnervation in human acute skin wounds treated by electrical stimulation. *Journal of investigative dermatology*, 137(3):737–747, 2017.

[78] Elaine Emmerson. Efficient healing takes some nerve: electrical stimulation enhances innervation in cutaneous human wounds. *Journal of Investigative Dermatology*, 137(3):543–545, 2017.

[79] Siwei Zhao, Abijeet Singh Mehta, and Min Zhao. Biomedical applications of electrical stimulation. *Cellular and molecular life sciences*, 77(14):2681–2699, 2020.

[80] IV Yannas, DS Tzeranis, and PTC So. Regeneration of injured skin and peripheral nerves requires control of wound contraction, not scar formation. *Wound Repair and Regeneration*, 25(2):177–191, 2017.

[81] Tingting Weng, Pan Wu, Wei Zhang, Yurong Zheng, Qiong Li, Ronghua Jin, Haojiao Chen, Chuangang You, Songxue Guo, Chunmao Han, and Xingang Wang. Regeneration of skin appendages and nerves: current status and further challenges. *Journal of Translational Medicine*, 18(1), 2 2020.

[82] Maja Schlereth, Daniel Stromer, Yash Mantri, Jason Tsujimoto, Katharina Breininger, Andreas Maier, Caesar Anderson, Pranav S. Garimella, and Jesse V. Jokerst. Initial investigations towards non-invasive monitoring of chronic wound healing using deep learning and ultrasound imaging. In *Informatik aktuell*, pages 261–266. Springer Fachmedien Wiesbaden, 2022.

[83] Abolfazl Zargari, Gerrald A. Lodewijk, Najmeh Mashhadi, Nathan Cook, Celine W. Neudorf, Kimiasadat Araghbidikashani, Stefany Rubio, Eva Hrabeta-Robinson, Angela N. Brooks, Lindsay Hinck, and S. Ali Shariati. Deepsea: An efficient deep learning model for single-cell segmentation and tracking of time-lapse microscopy images. 3 2021.

[84] Hy Dang. Wound healing modeling using partial differential equation and deep learning. 11 2021.

[85] Prashant Kumar, Swatantra Dhara, Ayan Gope, Jyotirmoy Chatterjee, and Subhamoy Mandal. Deep learning based skin-layer segmentation for characterizing cutaneous wounds from optical coherence tomography images. *45th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 06 2023.

[86] Gaetano Scebba, Jia Zhang, Sabrina Catanzaro, Carina Mihai, Oliver Distler, Martin Berli, and Walter Karlen. Detect-and-segment: A deep learning

approach to automate wound image segmentation. *Informatics in Medicine Unlocked*, 29:100884, 2022.

[87] Sebastian Wurthmann, Steffen Nägel, Eva Hadaschik, Swantje Schlott, Armin Scheffler, Christoph Kleinschnitz, and Dagny Holle. Impaired wound healing in a migraine patient as a possible side effect of calcitonin gene-related peptide receptor antibody treatment: A case report. *Cephalalgia*, 40(11):1255–1260, 6 2020.

[88] Alia Sadiq, Ahmed Shah, Marc Jeschke, Cassandra Belo, Muhammad Qasim Hayat, Sheeba Murad, and Saeid Amini-Nik. The role of serotonin during skin healing in post-thermal injury. *International Journal of Molecular Sciences*, 19(4):1034, 3 2018.

[89] Olaf Ronneberger. Invited talk: U-net convolutional networks for biomedical image segmentation. In *Informatik aktuell*, pages 3–3. Springer Berlin Heidelberg, 2017.

[90] Daniel B. Allan, Thomas Caswell, Nathan C. Keim, Casper M. van der Wel, and Ruben W. Verweij. soft-matter/trackpy: Trackpy v0.5.0, 2021.

[91] John D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science Engineering*, 9(3):90–95, 2007.

[92] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy,

Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Matplotlib: A 2d graphics environment. 2021.

[93] Ksenia Zlobina, Mohammad Jafari, Marco Rolandi, and Marcella Gomez. The role of machine learning in advancing precision medicine with feedback control. *Cell Report Physical Sciences*, 3:101149, 2022.

[94] Raquel Dias and Ali Torkamani. Artificial intelligence in clinical and genomic diagnostics. *Genome Medicine*, 11:70, 2019.

[95] C Krittanawong, H Zhang, Z Wang, M Aydar, and Kitai T. Artificial intelligence in precision cardiovascular medicine. *J Am Coll Cardiol*, 69:2657–2664, 2017.

[96] Dan M Roden and Alfred L George Jr. The genetic basis of variability in drug responses. *Nature reviews Drug discovery*, 1(1):37–44, 2002.

[97] Suzette J Bielinski, Janet E Olson, Jyotishman Pathak, Richard M Weinshilboum, Liewei Wang, Kelly J Lyke, Euijung Ryu, Paul V Targonski, Michael D Van Norstrand, Matthew A Hathcock, et al. Preemptive genotyping for personalized medicine: design of the right drug, right dose, right time—using genomic data to individualize treatment protocol. In *Mayo Clinic Proceedings*, volume 89, pages 25–33. Elsevier, 2014.

[98] Behnood Gholami, Wassim M Haddad, James M Bailey, and Allen R Tannenbaum. Optimal drug dosing control for intensive care unit sedation by using a hybrid deterministic–stochastic pharmacokinetic and pharmacodynamic model. *Optimal Control Applications and Methods*, 34(5):547–561, 2013.

[99] Regina Padmanabhan, Nader Meskin, and Wassim M Haddad. Closed-loop

control of anesthesia and mean arterial pressure using reinforcement learning. *Biomedical Signal Processing and Control*, 22:54–64, 2015.

[100] Ksenia Zlobina, Jiahao Xue, and Marcella Gomez. Effective spatio-temporal regimes for wound treatment by way of macrophage polarization: a mathematical model. *Frontiers in Applied Mathematics and Statistics*, 8:791064, 2022.

[101] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.

[102] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.

[103] Jean-Jacques E. Slotine. Sliding controller design for non-linear systems. *International Journal of Control*, 40:421–434, 1984.

[104] You Zheng, Jianxing Liu, Xinyi Liu, Dandan Fang, and L. Wu. Adaptive second-order sliding mode control design for a class of nonlinear systems with unknown input. *Mathematical Problems in Engineering*, 2015:1–7, 10 2015.

[105] Yung-Chiang Chung, Bor-Jiunn Wen, and Yu-Chen Lin. Optimal fuzzy sliding-mode control for bio-microfluidic manipulation. *Control Engineering Practice*, 2007.

[106] Héctor Carrión, Mohammad Jafari, Hsin-Ya Yang, Roslyn Rivkah Isseroff, Marco Rolandi, Marcella Gomez, and Narges Norouzi. Healnet-self-supervised acute wound heal-stage classification. In *International Workshop on Machine Learning in Medical Imaging*, pages 446–455. Springer, 2022.

[107] Ksenia Zlobina, Eric Malekos, Han Chen, and Marcella Gomez. Robust classification of wound healing stages in both mice and humans for acute and burn wounds based on transcriptomic data. *BMC Bioinformatics*, 24(1):166, Apr 2023.

[108] S al Guo and Luisa A DiPietro. Factors affecting wound healing. *Journal of dental research*, 89(3):219–229, 2010.

[109] Héctor Carrión, Mohammad Jafari, Michelle Dawn Bagood, Hsin-ya Yang, Roslyn Rivkah Isseroff, and Marcella Gomez. Automatic wound detection and size estimation using deep learning algorithms. *PLoS computational biology*, 18(3):e1009852, 2022.

[110] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a" siamese" time delay neural network. *Advances in neural information processing systems*, 6, 1993.

[111] Yuanwen Jiang, Artem A Trotsyuk, Simiao Niu, Dominic Henn, Kellen Chen, Chien-Chung Shih, Madelyn R Larson, Alana M Mermin-Bunnell, Smiti Mittal, Jian-Cheng Lai, et al. Wireless closed-loop smart bandage for chronic wound management and accelerated tissue regeneration. *bioRxiv*, pages 2022–01, 2022.

[112] Atte Kekonen, Mikael Bergelin, Max Johansson, Narender Kumar Joon, Johan Bobacka, and Jari Viik. Bioimpedance sensor array for long-term monitoring of wound healing from beneath the primary dressings and controlled formation of h2o2 using low-intensity direct current. *Sensors*, 19(11):2505, 2019.

[113] Henry C Lukaski and Micheal Moore. Bioelectrical impedance assessment of wound healing. *Journal of diabetes science and technology*, 6(1):209–212, 2012.

[114] Valentina Dini, Pietro Salvo, Agata Janowska, Fabio Di Francesco, Alessandro Barbini, and Marco Romanelli. Correlation between wound temperature obtained with an infrared camera and clinical wound bed score in venous leg ulcers. *Wounds: a compendium of clinical research and practice*, 27(10):274–278, 2015.

[115] Luther C Kloth. Electrical stimulation technologies for wound healing. *Advances in wound care*, 3(2):81–90, 2014.

[116] Min Zhao, Bing Song, Jin Pu, Teiji Wada, Brian Reid, Guangping Tai, Fei Wang, Aihua Guo, Petr Walczysko, Yu Gu, et al. Electrical signals control wound healing through phosphatidylinositol-3-oh kinase-γ and pten. *Nature*, 442(7101):457–460, 2006.

[117] Daniel J Cohen, W James Nelson, and Michel M Maharbiz. Galvanotactic control of collective cell migration in epithelial monolayers. *Nature materials*, 13(4):409–417, 2014.

[118] Cen Chen, Xue Bai, Yahui Ding, and In-Seop Lee. Electrical stimulation as a novel tool for regulating cell behavior in tissue engineering. *Biomaterials research*, 23:1–12, 2019.

[119] Chandan K Sen. Human wound and its burden: updated 2020 compendium of estimates. *Advances in wound care*, 10(5):281–292, 2021.

[120] Matthew Berezo, Joshua Budman, Daniel Deutscher, Cathy Thomas Hess, Kyle Smith, and Deanna Hayes. Predicting chronic wound healing time using machine learning. *Advances in wound care*, 11(6):281–296, 2022.

[121] John Selberg, Mohammad Jafari, Juanita Mathews, Manping Jia, Pattawong Pansodtee, Harika Dechiraju, Chunxiao Wu, Sergio Cordero, Alexander Flora, Nebyu Yonas, et al. Machine learning-driven bioelectronics for closed-loop control of cells. *Advanced Intelligent Systems*, 2(12):2000140, 2020.

[122] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.

[123] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256, 1992.

[124] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.

[125] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.

[126] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[127] Peter Lancaster and Leiba Rodman. *Algebraic riccati equations*. Clarendon press, 1995.

[128] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda

Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

[129] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

[130] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[131] James S Albus. Brains, behavior, and robotics. 1981.

[132] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[133] Andrey Bernstein, Yue Chen, Marcello Colombino, Emiliano Dall'Anese, Prashant Mehta, and Sean Meyn. Optimal rate of convergence for quasi-stochastic approximation. *arXiv preprint arXiv:1903.07228*, 2019.

[134] Shuhang Chen, Adithya Devraj, Andrey Bernstein, and Sean Meyn. Accelerating optimization and reinforcement learning with quasi-stochastic approximation. *arXiv preprint arXiv:2009.14431*, 2020.

[135] Zhifeng Zhang St'ephane Mallat. Matching pursuit with time-frequency dictionaries. In *IEEE Transactions on signal processing*. Citeseer, 1993.

[136] Alec Koppel, G Warnell, E Stump, P Stone, and Alejandro Ribeiro. Breaking bellman's curse of dimensionality: Efficient kernel gradient temporal difference. *arXiv preprint arXiv:1709.04221*, 2017.

[137] Alec Koppel, Garrett Warnell, Ethan Stump, and Alejandro Ribeiro. Parsimonious online learning with kernels via sparse projections in function space. *The Journal of Machine Learning Research*, 20(1):83–126, 2019.

[138] Jyrki Kivinen, Alexander J Smola, Robert C Williamson, et al. Online learning with kernels. In *NIPS*, pages 785–792, 2001.

[139] Amol Sasane. *Optimization in function spaces*. Courier Dover Publications, 2016.

[140] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006.

[141] Eurika Kaiser, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of koopman eigenfunctions for control. *Machine Learning: Science and Technology*, 2(3):035023, 2021.

[142] Katsuhiko Ogata. *Modern control engineering fifth edition*. 2010.

[143] Peter Van Overschee and BL0888 De Moor. *Subspace identification for linear systems: Theory—Implementation—Applications*. Springer Science & Business Media, 2012.

[144] Dimitri Bertsekas and John N Tsitsiklis. *Neuro-dynamic programming*. Athena Scientific, 1996.

[145] Yuanwen Jiang, Artem A Trotsyuk, Simiao Niu, Dominic Henn, Kellen Chen, Chien-Chung Shih, Madelyn R Larson, Alana M Mermin-Bunnell, Smiti Mittal, Jian-Cheng Lai, et al. Wireless, closed-loop, smart bandage with integrated sensors and stimulators for advanced wound care and accelerated healing. *Nature Biotechnology*, 41(5):652–662, 2023.

[146] Krister Järbrink, Gao Ni, Henrik Sönnergren, Artur Schmidtchen, Caroline Pang, Ram Bajpai, and Josip Car. The humanistic and economic burden of chronic wounds: a protocol for a systematic review. *Systematic reviews*, 6:1–7, 2017.

[147] Aline Costa de Oliveira, Daniel de Macêdo Rocha, Sandra Marina Gonçalves Bezerra, Elaine Maria Leite Rangel Andrade, Ana Maria Ribeiro dos Santos, and Lídya Tolstenko Nogueira. Quality of life of people with chronic wounds. *Acta Paulista de Enfermagem*, 32:194–201, 2019.

[148] Behnood Gholami, Wassim M Haddad, and James M Bailey. Ai in the icu: In the intensive care unit, artificial intelligence can keep watch. *IEEE Spectrum*, 55(10):31–35, 2018.

[149] Richard W Peck. Precision medicine is not just genomics: the right dose for every patient. *Annual review of pharmacology and toxicology*, 58:105–122, 2018.

[150] Ronald J Gordon. Standardized care versus precision medicine: Do we really need to wait for point-of-care testing? *Anesthesia & Analgesia*, 125(6):2161, 2017.

[151] Yin Long, Hao Wei, Jun Li, Guang Yao, Bo Yu, Dalong Ni, Angela LF Gibson, Xiaoli Lan, Yadong Jiang, Weibo Cai, et al. Effective wound healing enabled

by discrete alternative electric fields from wearable nanogenerators. *ACS nano*, 12(12):12533–12540, 2018.

[152] Norihiko Ohura, Ryota Mitsuno, Masanobu Sakisaka, Yuta Terabe, Yuki Morishige, Atsushi Uchiyama, Takumi Okoshi, Iizaka Shinji, and Akihiko Takushima. Convolutional neural networks for wound detection: the role of artificial intelligence in wound care. *Journal of Wound Care*, 28(Sup10):S13–S24, 2019.

[153] Andrew J Gale. Continuing education course# 2: current understanding of hemostasis. *Toxicologic pathology*, 39(1):273–280, 2011.

[154] MJ Portou, D Baker, D Abraham, and J Tsui. The innate immune system, toll-like receptors and dermal wound healing: A review. *Vascular pharmacology*, 71:31–36, 2015.

[155] Juliana Barr, Katayoun Zomorodi, Edward J Bertaccini, Steven L Shafer, and Eran Geller. A double-blind, randomized comparison of iv lorazepam versus midazolam for sedation of icu patients via a pharmacologic model. *The Journal of the American Society of Anesthesiologists*, 95(2):286–298, 2001.

[156] Thomas W Schnider, Charles F Minto, Pedro L Gambus, Corina Andresen, David B Goodale, Steven L Shafer, and Elizabeth J Youngs. The influence of method of administration and covariates on the pharmacokinetics of propofol in adult volunteers. *The Journal of the American Society of Anesthesiologists*, 88(5):1170–1182, 1998.

[157] Eiko Furutani, Keigo Tsuruoka, Shogo Kusudo, Gotaro Shirakami, and Kazuhiko Fukuda. A hypnosis and analgesia control system using a model

predictive controller in total intravenous anesthesia during day-case surgery. In *Proceedings of SICE Annual Conference 2010*, pages 223–226. IEEE, 2010.

[158] Kristian Soltesz, Jin-Oh Hahn, Tore Hägglund, Guy A Dumont, and J Mark Ansermino. Individualized closed-loop control of propofol anesthesia: a preliminary study. *Biomedical Signal Processing and Control*, 8(6):500–508, 2013.

[159] Jin-Oh Hahn, Guy A Dumont, and J Mark Ansermino. Robust closed-loop control of hypnosis with propofol using wavcns index as the controlled variable. *Biomedical Signal Processing and Control*, 7(5):517–524, 2012.

[160] James M Bailey and Wassim M Haddad. Drug dosing control in clinical pharmacology. *IEEE Control Systems Magazine*, 25(2):35–51, 2005.

[161] José D Martín-Guerrero, Faustino Gomez, Emilio Soria-Olivas, Jürgen Schmidhuber, Mónica Climente-Martí, and N Víctor Jiménez-Torres. A reinforcement learning approach for individualizing erythropoietin dosages in hemodialysis patients. *Expert Systems with Applications*, 36(6):9737–9742, 2009.

[162] Yufan Zhao, Donglin Zeng, Mark A Socinski, and Michael R Kosorok. Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics*, 67(4):1422–1433, 2011.

[163] Elena Daskalaki, Peter Diem, and Stavroula G Mougiakakou. Personalized tuning of a reinforcement learning control algorithm for glucose regulation. In *2013 35th Annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pages 3487–3490. IEEE, 2013.

[164] Regina Padmanabhan, Nader Meskin, and Wassim M Haddad. Reinforcement

learning-based control of drug dosing for cancer chemotherapy treatment. *Mathematical biosciences*, 293:11–20, 2017.

[165] Brett L Moore, Larry D Pyeatt, Vivekanand Kulkarni, Periklis Panousis, Kevin Padrez, and Anthony G Doufas. Reinforcement learning for closed-loop propofol anesthesia: a study in human volunteers. *The journal of machine learning research*, 15(1):655–696, 2014.

[166] Draguna L Vrabie, Kyriakos G Vamvoudakis, and Frank L Lewis. Optimal adaptive control and differential games by reinforcement learning principles. Institution of Electrical Engineers, 2013.

[167] Joshua L Proctor, Steven L Brunton, and J Nathan Kutz. Generalizing koopman theory to allow for inputs and control. *SIAM Journal on Applied Dynamical Systems*, 17(1):909–930, 2018.

[168] Dragoslav D Šiljak and AI Zečević. Control of large-scale systems: Beyond decentralized feedback. *Annual Reviews in Control*, 29(2):169–179, 2005.

[169] Athanasios C Antoulas. An overview of approximation methods for large-scale dynamical systems. *Annual reviews in Control*, 29(2):181–190, 2005.

[170] Paraskevas N Paraskevopoulos. *Modern control engineering*. CRC Press, 2017.

[171] Bernard O Koopman. Hamiltonian systems and transformation in hilbert space. *Proceedings of the National Academy of Sciences*, 17(5):315–318, 1931.

[172] Igor Mezić and Andrzej Banaszuk. Comparison of systems with complex behavior. *Physica D: Nonlinear Phenomena*, 197(1-2):101–133, 2004.

[173] Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41:309–325, 2005.

[174] Akhil Ahmed, Ehecatl Antonio del Rio-Chanona, and Mehmet Mercangöz. Learning linear representations of nonlinear dynamics using deep learning. *IFAC-PapersOnLine*, 55(12):162–169, 2022.

[175] Yingzhao Lian, Renzi Wang, and Colin N Jones. Koopman based data-driven predictive control. *arXiv preprint arXiv:2102.05122*, 2021.

[176] Regina Padmanabhan, Nader Meskin, and Wassim Haddad. Optimal adaptive control of drug dosing using integral reinforcement learning. *Mathematical biosciences*, 309:131–142, 2019.

[177] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

[178] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.

[179] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[180] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.

[181] Chiraz Atri, Fatma Z Guerfali, and Dhafer Laouini. Role of human macrophage polarization in inflammation during infectious diseases. *International journal of molecular sciences*, 19(6):1801, 2018.

[182] Fan Lu, Prashant G Mehta, Sean P Meyn, and Gergely Neu. Convex Q-learning. In *2021 American Control Conference (ACC)*, pages 4749–4756. IEEE, 2021.

[183] Fan Lu, Prashant Mehta, Sean Meyn, and Gergely Neu. Convex analytic theory for convex Q learning. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 4065–4071. IEEE, 2022.

[184] Fan Lu, G Prashant Mehta, P Sean Meyn, and Gergely Neu. Sufficient exploration for convex Q learning. *arXiv preprint arXiv:2210.09409*, 2022.

[185] Fan Lu and Sean Meyn. Convex Q learning in a stochastic environment: Extended version. *arXiv preprint arXiv:2309.05105*, 2023.

[186] Yanlin Zhou, Fan Lu, George Pu, Xiyao Ma, Runhan Sun, Hsi-Yuan Chen, and Xiaolin Li. Adaptive leader-follower formation control and obstacle avoidance via deep reinforcement learning. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4273–4280. IEEE, 2019.

[187] Constance Duchesne, Nadira Frescaline, Jean-Jacques Lataillade, and Antoine Rousseau. Comparative study between direct and indirect treatment with cold atmospheric plasma on in vitro and in vivo models of wound healing. *Plasma Medicine*, 8(4):379–401, 2018.

[188] P.W.D. Charles. Project title. `https://github.com/charlespwd/project-title`, 2013.