Controllable Synthetic Data Generation for Cutaneous Ulcer Tissue Segmentation

Júlia Garcia Torné

Abstract

Deep learning approaches for medical image segmentation require large amounts of annotated data, yet obtaining high-quality labels for wound and ulcer analysis is expensive, time-consuming, and clinically challenging. In cutaneous ulcer tissue segmentation, datasets are often limited, imbalanced, and highly variable in terms of illumination, tissue appearance, and acquisition conditions. These limitations reduce the generalization capability of segmentation models, particularly for underrepresented tissue classes.

This project proposes a controllable synthetic data generation framework to create realistic ulcer images conditioned on semantic tissue masks. The objective is to generate coherent image–mask pairs that preserve both anatomical plausibility and tissue-level consistency while increasing dataset diversity. The proposed approach will explore diffusion-based generative models, leveraging conditional image synthesis techniques such as ControlNet, latent diffusion, and parameter-efficient fine-tuning strategies including LoRA.

The workflow will consist of three main stages. First, semantic masks representing different tissue regions will be used as structural conditioning inputs. Second, a generative model will synthesize realistic ulcer images aligned with the mask geometry and tissue distribution. Finally, the generated synthetic samples will be integrated into downstream tissue segmentation pipelines to evaluate their impact on segmentation performance.

The project will analyze image realism, structural consistency between generated images and masks, and the usefulness of synthetic samples for improving segmentation robustness. Performance will be evaluated using both qualitative inspection and quantitative segmentation metrics on real validation data.

The expected contribution of this work is a scalable synthetic data generation pipeline capable of alleviating annotation scarcity in medical imaging while improving the robustness and generalization of ulcer tissue segmentation systems.

1 Introduction

Semantic segmentation of chronic wounds and cutaneous ulcers plays a key role in clinical assessment, tissue monitoring, and treatment planning. However, training robust segmentation models requires large annotated datasets, which are difficult to obtain due to the need for expert labeling and patient privacy constraints.

Synthetic data generation has recently emerged as a promising strategy to address data scarcity in medical imaging. Generative models can increase dataset diversity, balance minority classes, and improve the robustness of downstream tasks. In particular, diffusion models have demonstrated state-of-the-art capabilities in controllable image synthesis.

This project focuses on generating synthetic ulcer images conditioned on semantic tissue masks. By controlling tissue layout and structure, the generated samples can provide realistic and clinically meaningful image–mask pairs suitable for training segmentation networks.

2 Objectives

The main objectives of this project are:

- Develop a controllable synthetic image generation pipeline for cutaneous ulcers.
- Generate realistic ulcer images conditioned on semantic tissue masks.
- Study diffusion-based conditioning methods for medical image synthesis.
- Evaluate the usefulness of synthetic data for improving tissue segmentation performance.
- Analyze the realism and structural consistency of generated samples.

3 Methodology

The proposed methodology consists of the following stages:

3.1 Dataset Preparation

Real ulcer datasets containing RGB images and semantic tissue masks will be preprocessed and standardized. Tissue classes may include granulation, slough, necrotic tissue, and surrounding skin.

3.2 Conditional Synthetic Generation

Diffusion-based generative models will be explored to synthesize ulcer images conditioned on segmentation masks. The project may investigate:

- Latent Diffusion Models (LDMs)
- ControlNet conditioning
- LoRA fine-tuning
- Inpainting-based refinement
- Style conditioning from real ulcer images

The goal is to preserve structural correspondence between masks and generated images while maintaining visual realism.

3.3 Segmentation Evaluation

Synthetic samples will be incorporated into segmentation model training pipelines. The effect of synthetic augmentation will be evaluated using metrics such as:

- Dice Similarity Coefficient (DSC)
- Intersection over Union (IoU)
- Pixel Accuracy

Performance will be compared against models trained exclusively on real data.

4 Expected Outcomes

The project is expected to produce:

- A controllable synthetic ulcer image generation framework.
- A synthetic image–mask dataset for tissue segmentation.
- Experimental analysis of synthetic data utility in medical segmentation tasks.
- Insights into diffusion-based conditioning methods for biomedical image generation.

5 Conclusion

This work aims to address the challenge of limited annotated medical datasets through controllable synthetic data generation. By leveraging diffusion-based models conditioned on semantic tissue masks, the project seeks to create realistic and structurally consistent ulcer images that improve downstream segmentation performance and robustness.