NDSU NORTH DAKOTA STATE UNIVERSITY

National Science Foundation

DIABETIC FOOT ULCER SEGMENTATION USING MACHINE LEARNING

BY: CARLOS ANDREWS, DR. SIMONE LUDWIG, SHIVANI GAUTAM

UNIVERSIDAD DE CHILE

INTRODUCTION

In 2021, diabetes affected over 500 million people and was the leading cause of nontraumatic lower extremity amputations. Diabetic foot ulcers (DFUs) affect 15% of diabetic patients, with 6% hospitalized for complications. In the U.S., 14-24% of patients with foot ulcers face amputation. Early DFU detection can reduce amputation risk, but it is challenging [1] [2][3]. Machine learning models, especially image segmentation, help identify DFUs in images. The 2022 DFU Challenge winner, HarDNet-DFUS, achieved a dice score of 0.7287 [6]. We have implemented this, as well as the Eff-UNet, a compound scaling method applied to the UNet architecture.

Figure 1. DFU segmentation predicted by Eff-UNet

BACKGROUND

The HarDNet (Harmonic Dense Network) is designed to reduce memory traffic during inference, based on Dense-Net but optimizing memory use by flushing unnecessary layers. The HarDNet-MSEG, a variant tailored for medical image segmentation, led to the HarDNet-DFUS model. HarDNet-DFUS, with a dice score of 0.7287, won the 2022 Diabetic Foot Ulcer Challenge [10][11][12].

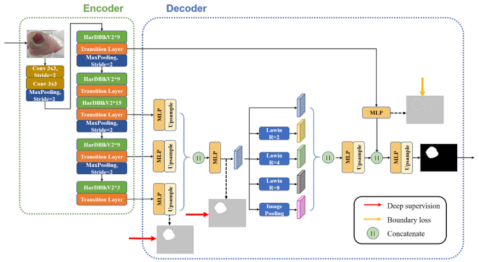

Figure 2. HarDNet-DFUS Architecture

Conventional scaling of convolutional neural networks (CNNs) often involves random adjustments to depth, width, or resolution, requiring extensive manual tuning. EfficientNet addresses this by using compound scaling, uniformly scaling width, depth, and resolution with fixed ratios [16]. EfficientNet experiments with MobileNets and ResNets. Eff-UNet integrates EfficientNetB7 into U-Net for improved efficiency [17].

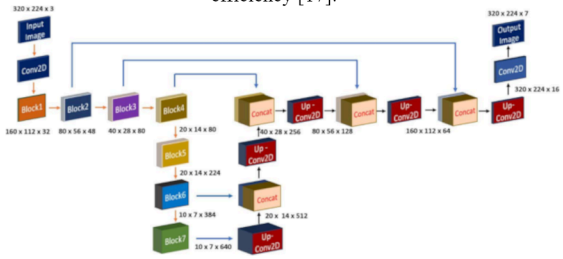

Figure 3. Eff-UNet Architecture

METHOD

The study used CCAST resources at North Dakota State University. Both models were trained and tested on the 2022 DFUC dataset. HarDNet-DFUS followed its GitHub guidelines, while Eff-UNet was implemented using TensorFlow with EfficientNetB7 and specific preprocessing techniques.

RESULTS

Training HarDNet-DFUS yielded a best dice score of 0.8234, slightly higher than the original mean of 0.7287 due to preprocessing variations. The Eff-UNet, trained for 30 epochs with a 3:2 Dice to BCE loss ratio, achieved a mean dice score of 0.65, indicating potential for improvement.

CONCULSION

The HarDNet-DFUS remains a state-of-the-art model, showing improved results over previous models. Reproduced results suggest it might outperform earlier claims. The Eff-UNet, while not surpassing existing models, provides adequate predictions and could benefit from further research and modifications. This study advances machine learning in healthcare, highlighting effective models and encouraging ongoing research to enhance medical applications. Future work may adapt these methods to achieve better results, contributing to the progress in healthcare machine learning.

References