

Generating 3D Bio-Printable Patches Using Wound Segmentation and Reconstruction to Treat Diabetic Foot Ulcers

Han Joo Chae Seunghwan Lee Hyewon Son Seungyeob Han Taebin Lim
 ROKIT Healthcare, Inc.

{hanjoo.chae, seunghwan.lee, hyewon.son, seungyeob.han, taebin.lim}@rokit.co.kr




Figure 1. The procedure of the DFU treatment includes (a) wound image capturing, (b) hybrid 3D patch generation combining 2D semantic segmentation and 3D reconstruction, (c) printing the patch using a 3D bio-printer, and (d) applying the printed patch to the wound.

Abstract

We introduce AiD Regen, a novel system that generates 3D wound models combining 2D semantic segmentation with 3D reconstruction so that they can be printed via 3D bio-printers during the surgery to treat diabetic foot ulcers (DFUs). AiD Regen seamlessly binds the full pipeline, which includes RGB-D image capturing, semantic segmentation, boundary-guided point-cloud processing, 3D model reconstruction, and 3D printable G-code generation, into a single system that can be used out of the box. We developed a multi-stage data preprocessing method to handle small and unbalanced DFU image datasets. AiD Regen’s human-in-the-loop machine learning interface enables clinicians to not only create 3D regenerative patches with just a few touch interactions but also customize and confirm wound boundaries. As evidenced by our experiments, our model outperforms prior wound segmentation models and our reconstruction algorithm is capable of generating 3D wound models with compelling accuracy. We further conducted a case study on a real DFU patient and demonstrated the effectiveness of AiD Regen in treating DFU wounds.

1. Introduction

Diabetic foot ulcer (DFU) is the most costly and devastating complication of diabetes mellitus, which affects 15% of diabetic patients during their lifetimes [61]. Poor

blood circulation and high glucose levels decrease diabetic patients’ ability to heal, and every 20 seconds, a leg is lost through amputation due to infection and necrosis [3]. With the urgent need for a novel treatment method, the integration of 3D printing and bio-materials is paving the way toward devising innovative solutions in regenerative medicine [48]. Kesavan et al. [35] have applied 3D bio-printed patches, that consist of minimally manipulated extracellular matrix (MA-ECM) harvested from autologous fat tissues, to DFU wounds and shown a complete healing with re-epithelialization only within about four weeks.

However, such use of 3D bio-printing technology is not accessible for many medical practitioners as it requires a complex process of 3D scanning and modeling. Conventional medical scanning methods, such as CT and MRI scans, often demand large, special devices and environments as well as experts who can operate them. In addition, the results of these scans usually include unnecessary parts and peripherals so that hospitals need to hire CAD engineers for post-processing the results or rely on outsourcing; either way, cost increases are inevitable. Also, since treatments of skin wounds like DFU are usually followed by debridement, a surgical removal of damaged tissue, the 3D patch creation must be done as quickly as possible to prevent debrided wounds from contamination. External post-processing and outsourcing might not be the best solutions in such time sensitive situations. Though portable industrial 3D scanners [53] lowered spatial and technical barriers, they still require users to understand correct scanning

methods, and the post-processing needs remain the same.

3D object reconstruction using RGB-D images [65] has been actively studied to automate modeling tasks; segmenting each part of an object can be done in real-time [42]. Unfortunately, however, most existing wound datasets consist only of 2D images, and available medical data are often small or fragmented due to their sensitive nature; that is, preparing a large RGB-D or 3D wound training dataset might be difficult. Hybrid approaches, that combine features extracted from 2D images with depth data, could be good solutions to this as they have shown convincing performance in reconstructing wounds [19, 21, 64].

In this work, we introduce *AiD Regen*, a novel system that generates 3D models based on the shape and geometry of wounds combining 2D image segmentation and 3D reconstruction so that they can be printed via 3D bioprinters during the surgery to treat DFUs. *AiD Regen* seamlessly binds the full pipeline, which includes RGB-D image capturing, semantic segmentation, boundary-guided point-cloud processing (BGPCP), 3D model reconstruction, and 3D bio-printable G-code generation, into a single system that can be used out of the box. We also develop a multi-stage data pipeline to handle small and unbalanced characteristics of a DFU image dataset. Reflecting requirements of medical practitioners, *AiD Regen*'s human-in-the-loop (HITL) machine learning (ML) interface enables clinicians to not only create 3D regenerative patches with just a few number of touch interactions, but also customize and confirm wound boundaries. We evaluate the performance of our 2D segmentation and 3D model generation methods through a series of experiments and demonstrate the effectiveness of *AiD Regen* by conducting a case study on a real DFU patient.

2. Related Work

2.1. 3D Reconstruction and Segmentation

Along with the advancements in RGB-D cameras, reconstructing real-world environments to 3D models has been actively studied [65]. In addition, efforts have been exerted to separate each object in the given scene as in RGB-D based semantic segmentation [14, 25, 32, 38, 40, 41]. However, these research using public 3D semantic segmentation dataset are yet immature and their performances are reported at below 50 % mIoU [27]. Moreover, gathering 3D segmentation data is difficult since it requires high quality sensors as well as 3D domain expertise.

In terms of medical applications, hybrid approaches, that combine 2D image features with depth data to reconstruct 3D objects, have been most commonly used. Filko et al. [19] reconstructed a 3D scene using the marching cubes algorithm and segmented wounds using color histograms and k-Nearest-Neighbor algorithm. Gutierrez et al. [21]

used deep learning methods to perform semantic segmentation on DFU wounds and merged the results with the 3D foot model reconstructed from a sequence of RGB images. Our *AiD Regen* system also takes a hybrid approach to exploit well-established, deep-learning-based 2D semantic segmentation to find accurate boundaries of DFU wounds and use them to create 3D regenerative patches.

2.2. 2D Semantic Segmentation

Since fully convolutional network has been proposed [44], deep learning approaches based on convolutional neural network (CNN) have shown great successes in semantic segmentation. Spatial pyramid pooling [26, 63] has shown its effectiveness in obtaining details of objects, and DeepLab [12] has been one of the most successful networks utilizing this technique. DeepLab applies atrous convolution to get multi-scale features and then concatenate the resulting features; this helps with getting feature maps from various receptive fields while keeping the high computational efficiency.

Attention model [6, 11, 31] also has been successful in semantic segmentation as it weighs the importance of scale or channel. SENet [31] utilized channel attention to adaptively recalibrate channel-wise feature responses based on inter-dependencies between channels. Chen et al. [11] utilized scale attention to focus on weighting multi-scale features.

Transformer models are also actively being applied to semantic segmentation tasks [7, 43]. These research use self-attention mechanisms rather than CNNs as their main structures. While the Transformer models are competent in semantic segmentation, they perform worse than CNN-based networks when the training dataset is not large enough [16].

In contrast to general objects, relatively little attention has been paid towards wound segmentation. Wang et al. [58] systematically compared VGG16, SegNet, U-Net, Mask-RCNN, and MobileNetV2 networks on wound segmentation tasks and reported that MobileNetV2 showed the best performance. While there is yet no large public dataset specifically designed for DFU segmentation, a few models have been proposed [20, 45, 49] and CNN-based models have shown the most reliable performances due to their data-efficient nature.

In this paper, we build a new DFU dataset (RKT-DFU-800) consisting of real-world DFU wound images taken by clinicians during surgeries or examinations and overcome the nature of unbalanced real-world data with wound-specific augmentation techniques. Inspired by the work of Azad et al. [6], we apply the attention mechanism to DeepLabv3 to preserve details of wound boundaries.

2.3. Medical Machine Learning Systems

While a proliferation of studies on individual ML tasks, such as semantic segmentation [12, 13, 54, 62] and 3D re-

construction [23, 33], has been published each year, relatively little studies have focused on applying and adapting those techniques to specific real-world problems, which require significant algorithmic and engineering work [34, 57]. In fact, only a few medical areas, such as prosthesis [30] and dentistry [15], tried to bind the process of 3D scanning, modeling, and printing. However, they all involved the use of CAD software which requires a certain level of expertise. Our goal is to not only build an automatic system that encompasses the full pipeline from scanning skin wounds to generating 3D-printable patches, but also provide a user-friendly interface that does not require special skills.

In addition, any ML system can fail especially while running in the wild; Bengio et al. [8] have reported that the performance of today’s best ML systems often break down when they go out from the lab. For instance, some of the most popular and large image datasets possess geometric bias so that the ML models trained with those data resulted in errors with samples that are not included in the dataset [24, 52]. Such errors can be disastrous in risk-sensitive fields like medical systems [60]. To minimize the dataset bias, we collected real world DFU data from multiple medical centers globally.

Furthermore, the HITL ML approach has been proposed in multiple studies of medical image analysis [9] to maximize the performance while reducing the risk of malfunctioning. Amrehn et al. [4] utilized masks obtained from user scribbles to perform hepatic lesion segmentation and allowed users to add seed points if they were not satisfied with the prediction. Wang et al. [59] proposed the image-specific fine-tuning method based on user scribbles. We adopt HITL to our user interface to reduce risks from unexpected predictions by allowing clinicians to modify and confirm the results. Also, our user-friendly interface and interaction do not demand separate learning process.

3. Method

3.1. System Overview

The AiD Regen system encompasses the full-pipeline from capturing RGB-D images to generating 3D bio-printable patches to treat patients with DFUs. AiD Regen only asks users to take a single photo and confirm the suggested wound boundaries without requiring technical knowledge about 3D modeling, and the rest of the complex patch generation process is taken care of by the system. The entire pipeline of the system is depicted in Fig. 2. When a clinician takes a photo with an RGB-D camera, the wound segmentation module analyzes the image and returns the detected wound boundary along with the segmentation map. The clinician can refine the result using the HITL-incorporated user interface. Then, the depth map and the refined boundary information are sent to the server.

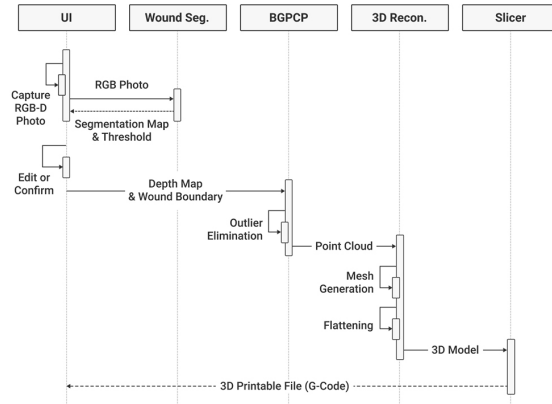


Figure 2. A sequence diagram of AiD Regen. Users can interact with the system via the user interface module and the rest of the modules seamlessly take care of the 3D-printable patch generation process.

To increase the robustness against sensor noises near rough edges of the wound [36], the boundary-guided point cloud processing (BGPCP) module eliminates outliers along the given boundary. The 3D reconstruction module generates a 3D surface mesh with the processed point clouds and then flattens it out to reduce printing time and cost. The final 3D model is created by using the flattened mesh and user-given thickness and fed to the slicer module to be converted into printing instructions known as G-code. By sending this G-code to a specified 3D bio-printer, the clinician can print out regenerative patches and treat the patients with DFUs.

3.2. Capturing 3D Wound Images

While various types of 3D cameras exist, the structured light sensors are known to have the highest precision when taken from a short distance [1]. As the TrueDepth camera on Apple devices uses the structured light sensors and is reported to have better depth accuracy compared to other LiDAR or stereo-vision based sensors [1, 56], AiD Regen utilizes the 12-megapixel front-facing camera of iPad Pro 11 (second generation) for capturing RGB-D wound images. The depth map is processed with temporal interpolation to smooth noises and remove holes and the resulting image consists of 2316×3088 RGB image and 480×640 depth map. Since taking a photo of a patient’s wound using the front-facing camera is difficult, we mounted a silver-coated first surface mirror accessory onto the device (Fig. 1a).

3.3. Wound Image Segmentation

Our first step towards training a deep-learning wound segmentation model was creating a DFU image dataset (RKT-DFU-800). To reduce the dataset bias as much as possible, we collected data from institutional review boards (IRB) approved clinical trials conducted globally in four




Figure 3. Our proposed model adopts the attention mechanism to DeepLabv3 to preserve details of wound boundaries.

different medical centers: Foot and Ankle Specialists of Southwest Virginia (Virginia, USA), Korea University Guro Hospital (Seoul, Korea), Ankara City Hospital (Ankara, Turkey), and Hycare Super Specialty Hospital (Tamil Nadu, India). Photos were taken without any fixed conditions to reflect real-world uses. RKT-DFU-800 consists of 818 unique DFU images from 63 patients. The number of images per patient varies from 1 to 68 and the image size varies from 143×176 to 4128×3096 . The images of the dataset not only contain DFU wounds but also other complications such as bleeding, slough, and necrosis. To prepare the dataset for the semantic segmentation task, two DFU experts independently labeled the data and then discussed until they reached a consensus. The dataset only contained foot focused images without any individual identifiers.

As we did not control the photo capturing process in order to reflect real-world scenarios, our dataset contains three different types of imbalances. First, DFU areas are sparse in images; only 3.4% of the total pixels are classified as DFU and the rest as background. In this case, the majority background class dominates the net gradient, and thus, without careful data engineering, the imbalance can constrain the model from learning features of the minority DFU class which might be more important [5]. Second, the number of images per patient highly varies; about half of the images are from 11 out of 67 patients. Since feature extractors are likely to focus more on the wounds of the patient with greater number of images, generalizing a deep learning model might be difficult. Third, relatively small number of bleeding samples are included in the dataset. While DFU wounds can have many different visual characteristics like dark necrosis tissues, we found bleeding as one of the most critical complications of the wound since bleeding near the wound boundary can confuse the model from determining accurate boundaries.

To handle the imbalances, we put several data augmentation stages as preprocessing. As the first stage, we resized input images in a way that the shorter edge of each image has 800 pixels. Next, to balance the number of images per patient, we applied a deterministic oversampling

method following Eq. (1):

$$S_{N_t} = \max(\min(C, N_M) - N_t, 0) \quad (1)$$

where N_M is the largest number of images that a single patient has in the dataset and N_t is the current number of images of patient t . C is a clamping parameter that sets the maximum number of oversampling images. S_{N_t} is, finally, the number of images to oversample for patient t . We set the clamping parameter C to 20.

As the third stage, we performed another oversampling to balance bleeding and non-bleeding wound images. In this stage, we manually scanned through the dataset to identify bleeding samples and empirically determined the amount of oversampling; we analyzed the validation losses and chose the best value. For all the oversampled images during the second and third stages, we applied random rotation, shear transformation, and scaling; the random rotation was ranged between 30° and 330° , the shear transformation matrix $\begin{pmatrix} 1 & s_x \\ s_y & 1 \end{pmatrix}$ spanned between -0.25 and 0.25 , and the scale factor ranged from 0.5 to 2.

As the fourth stage, we applied eight randomized augmentation techniques: adding Gaussian noise, blurring, swapping color channels, luminance modification, flipping, rotating, scaling, and applying shear transformation. As a result, eight additional images were created per image. This stage was performed before the training to efficiently manage the dataset versions and reproducibility of experiments.

The last stage was designed to handle the sparseness of DFU areas in the images; only 3.4% of the total pixels are classified as DFU. In this stage, we randomly sampled two wound pixels and six background pixels from each image and took 256×256 patches centered around those pixels, guaranteeing a quarter of the patches to be wound samples. Different from the previous stages, these patches were generated every iteration during the training.

We used a CNN-based model, DeepLabv3 [10], and applied the attention mechanism to it. The data-efficient characteristics of CNN-based models have shown their advantages in wound segmentation tasks using small datasets [37, 47, 58]. The architecture of our model is presented in Fig. 3.

From atrous spatial pyramid pooling layer, we could obtain features reflecting different fields of view. Each feature goes through two parallel and identical networks, which consist of 3×3 convolution layer and 1×1 convolution layer in series, to calculate score maps and attention maps using sigmoid and Softmax activations, respectively. The resulting score maps represent the features of the corresponding fields of view, and the attention maps show which field of view is more important at each pixel. The final segmentation map is obtained by upsampling the attention-weighted sum of the score maps.

For training, we split our dataset into training (646 images), validation (81 images), and test (91 images) sets. We constrained images of a patient included in one set to not appear in the other sets. We pretrained our model using the subset of COCO 2017 [39] and VOC2012 [17]. Images in the training set was passed through our data preprocessing and the total of 15,255 images were obtained. We set the batch size to 32 patches from four different images and the model was trained using four V100 GPUs. We stopped the training when the validation loss did not decrease for 1,000 iterations. The loss function was binary cross entropy and the stochastic gradient descent was used as an optimizer with learning rate of 6.8×10^{-4} without scheduling. For regularization, drop rate was set to 0.4, and L2 regularization coefficient was set to 2×10^{-4} . On test time, our model takes a 2D image as an input and returns the segmentation map and the default threshold value so that the users can review the result using our interactive user interface.

3.4. Human in the Loop

No ML models are perfect and we have to admit that our model might occasionally produce unexpected predictions; however, such errors can be disastrous in the risk-sensitive medical field. Moreover, regardless of the accuracy of the prediction, clinical decisions can vary from clinician to clinician [55]. To maximize the performance of our model and handle inter-operator variability while preventing system malfunction, we adopted the HITL design to our user interface (Fig. 4). It consists of three major interactions: i) seed point declaration, ii) adaptive thresholding, and iii) manual drawing or modification.

Firstly, when a user declares a seed point by touching a point within the wound area on the image, AiD Regen displays a single wound area that includes the given seed point. Starting from a seed point is useful especially when multiple wounds exist in a single image and the clinician needs to confirm the intended target. When the deep learning model outputs a segmentation map, it usually applies fixed thresholds to decide the final classification. However, depending on the image, a different threshold value could give better result. AiD Regen provides a slider so that the user can dynamically change the threshold while visually monitoring

the changes in the segmentation result. Since the segmentation scores change gradually near the wound boundaries, when visualized, changing the threshold has the practical effect of shrinking or expanding the boundary (Fig. 4b). While the user can simply accept the provided segmentation result without any additional interaction, AiD Regen allows them to completely redraw the wound boundary or modify the given boundary by moving key points (Fig. 4c). Since all the drawing or modifications are done on a 2D image, our interactions are significantly easier to use compared to existing modeling solutions which require the use of complicated CAD tools.

3.5. 3D Model Generation

When the BGPCP module receives the wound boundary confirmed by the clinician, it takes the pixels that are within the given boundary to convert them into a point cloud. While we have obtained an accurate boundary via semantic segmentation and user interaction, depth noises near the edges are still present because the sensors can get confused when a point lies right in between the background and the wound [36]. Here, we suggest a BGPCP method that efficiently performs the radius based noise filtering [22] by limiting the processing area around the given wound boundary; we compute the distance transform [18] on the boundary and generate a thin processing area by setting the distance parameter to $0.1 \times \sqrt{A_{boundary}}$ where $A_{boundary}$ is the area within the boundary (Fig. 5). The BGPCP method reduces the computation time of point cloud filtering by limiting the calculation area and is especially effective when the target wound is large.

Then, using the pixels that correspond to the points in the processed point cloud, we generate 2D mesh surfaces. We first create a convex-hull mesh using Delaunay triangulation and refine it to reflect the concavity of the wound area by removing all the edges whose midpoints lie outside of the boundary. The final 2D mesh is converted into a 3D mesh using the intrinsic matrix and depth map.




Figure 4. Our HITL-incorporated user interface. (a) AiD Regen suggests a wound boundary at the user-provided seed point. (b) The boundary shrinks when the user moves the threshold slider upwards. (c) The user can finetune the boundary by moving each point or manually redraw the boundary in necessary.

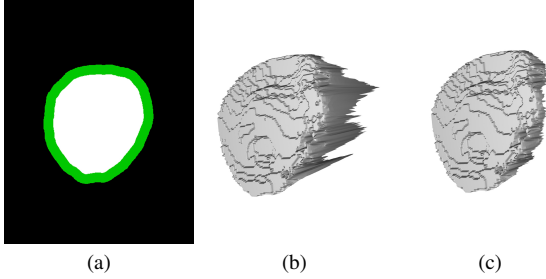


Figure 5. An example of BGPCP. (a) Green processing area is displayed on the segmentation map. (b) The original surface mesh with noise. (c) The BGPCP processed surface mesh.

3.6. Flattening and Slicing

Flattening complex 3D meshes is critical in 3D bio-printing since taller models need higher supporters, consuming greater amount of filament and bioink. Exploiting the flexible characteristics of 3D regenerative patches made with MA-ECM, we flatten the mesh surface using the boundary-first flattening (BFF) algorithm [51]. Lastly, using the flattened mesh surfaces, we create a 3D patch model with the user-given thickness and export it as an STL file.

Commercial 3D printers usually have their own slicer software that converts 3D models into a printable format. In order to automate the entire pipeline, we ported a slicer, NewCreatorK [29], that is intended to be used with a 3D bio-printer, Dr. INVIVO [28], to a cloud server connected with our gateway so that it can take an STL model as an input and return a G-code file as the output. Once the generated G-code is downloaded to the local device running AiD Regen, the user can send it over to the bio-printer to physically print out the MA-ECM patch.

4. Experiments

4.1. Segmentation Performance Comparison

We evaluated our model on the public wound dataset [58] and compared the results against the previous models; since there was no performance report using DeepLabv3 on the

Model	Recall	Precision	F1-Score
VGG16 [58]	0.7835	0.8391	0.8103
SegNet [58]	0.8649	0.8649	0.8503
U-Net [58]	0.9129	0.8904	0.9015
Mask-RCNN [58]	0.8640	0.9430	0.9020
MobileNetV2 [58]	0.8976	0.9086	0.9030
MobileNetV2+CCL [58]	0.8997	0.9101	0.9047
DeepLabV3 [10]	0.9207	0.9191	0.9199
Our model	0.9287	0.9172	0.9229

Table 1. Performance comparisons on the public wound dataset [58]. Our model outperforms the previous models in recall and F1-score.

Model	S2	S3	HITL	mIoU	Recall	Precision	F1-Score
U-Net	✓	✓		0.7770	0.8733	0.8757	0.8745
MobilenetV2	✓	✓		0.8494	0.9242	0.9131	0.9186
DeepLabV3	✓	✓		0.8707	0.9342	0.9277	0.9309
Our model				0.5944	0.6153	0.9460	0.7456
Our model	✓			0.6867	0.7075	0.9590	0.8143
Our model		✓		0.8753	0.9286	0.9384	0.9335
Our model	✓	✓		0.8840	0.9387	0.9382	0.9385
Our model	✓	✓	✓	0.9043	0.9517	0.9478	0.9497

Table 2. Results of our data processing and HITL interface on the validation set. S2 denotes the stage 2 augmentation which balances the number of images per patient, S3 denotes the stage 3 augmentation which oversamples the bleeding images, and HITL denotes that the HITL was simulated.

same dataset, we also trained and evaluated DeepLabv3 for comparison. We firstly resized the training set to 800×800 . We only applied the fourth stage of our data preprocessing as an augmentation without any oversampling or patch sampling. In the test step, the resulting segmentation map was downsized to 224×224 to match the original size of the label. As shown in Tab. 1, our model outperforms the previous models on recall (0.93) and F1-score (0.92).

We also evaluated the performance of our model using DFU-specific RKT-DFU-800 dataset. We trained U-Net [50], MobilenetV2 [58], DeepLabv3 [10], and our model by equally applying our data preprocessing and compared the performances on the validation set. As shown in Tab. 2, our model outperformed all the other models on mIoU (0.88) and F1-Score (0.94).

Furthermore, to evaluate the effectiveness of our data preprocessing, we tested our model using the same dataset but omitting one or both of the stages two and three. As shown in Tab. 2, the model performed the worst with mIoU of 0.59 when both stages were omitted and the best with mIoU of 0.88 when all the data preprocessing was applied. Moreover, to demonstrate the maximum possible performance using the HITL interface, we simulated the best use by setting the optimal segmentation threshold for each image individually. We found that by using our HITL interface, the performance can be raised up to 0.90 mIoU.

Lastly, we report the performance of our model on the test set in Tab. 3. The HITL interface increased the overall performance of mIoU, recall, precision, and F1-score.

4.2. 3D Model Accuracy

To measure the accuracy of 3D reconstructed wound model, we used an artificial foot model [46], which has

HITL	mIoU	Recall	Precision	F1-Score
	0.8246	0.9465	0.8650	0.9039
✓	0.8795	0.9329	0.9388	0.9359

Table 3. The performances of our model on the test set when the HITL was simulated or not.




Figure 6. An artificial foot model with DFU wounds and the ground truth surface meshes.

three different DFU wounds on the sole. As the ground truth, we scanned the model using an industrial 3D scanner, EinScan Pro 2X [53], and measured the surface area of each wound (Fig. 6). As an experiment, we captured RGB-D wound photos from a distance of 20 cm at varying angles from -30° to 30° in 10° increments. Five photos were taken from each angle and to focus solely on the reconstruction accuracy, we manually labeled each image to match the 3D scanned ground truth and compared the surface areas measured by AiD Regen against the ground truth. The results show that AiD Regen’s reconstruction accuracy ranged from 93.83 % to 97.07 % per wound types and from 94.75 % to 96.09 % per angles (Tab. 4). The total mean absolute error was $\mu = 0.92 \text{ cm}^2$, $\sigma = 1.05 \text{ cm}^2$ and the total average accuracy was 95.27 %.

4.3. Case Study

We conducted a case study to demonstrate the effectiveness of AiD Regen during a real clinical trial at the Catholic University of Korea Eunpyeong St. Mary’s Hospital (Seoul, Korea. IRB PC21EOSE007). The patient was 67-year-old woman with Type 2 diabetes. A doctor generated a 3D patch model using AiD Regen and printed an MA-ECM patch using a 3D bio-printer, Dr. INVIVO. As shown in Fig. 7a, the patch generated by AiD Regen was 97.97 % identical in area to the one created by an expert using a CAD modeling tool. The patient’s wound was healed drastically as 96 % of the wound was closed within just 34 days after the treatment (Fig. 7b). In addition, during each follow-up visit, we measured the wound size using both AiD Regen and the conventional method which involves a tape measure and ImageJ [2] software (Fig. 8). The measurements of AiD Regen and ImageJ were similar, both indicating the decreasing trend on every visit. AiD Regen’s measurements were on average 0.30 cm^2 greater than those of ImageJ.

4.4. Discussion

In order to build a robust ML model using small and imbalanced dataset, we introduced the multi-stage data pre-




Figure 7. (a) The AiD Regen generated patch (top) and the expert created patch using a CAD modeling tool (bottom) were very similar with only 2.03% difference in volume, and (b) 96% of the wound was healed within 34 days after the treatment.




Figure 8. The wound area measurements using AiD Regen and ImageJ for every follow-up visit.

processing. The results in Tab. 2 show that our data preprocessing successfully overcame such imbalances as our segmentation model achieved 0.8840 mIoU when the full data preprocessing was applied; this is 29% better than the one without applying stage 2 and 3 augmentations. In addition, our model, adopting the attention mechanism to DeepLabv3 network in order to preserve details of the wound boundaries, outperformed previous models on both the public wound dataset [58] as well as RKT-DFU-800.

Our HITL approach plays an important role in managing the segmentation performance as well as inter-operator variability; the boundary judgement and the treatment strategy can vary depending on the clinician. Our simple but highly customizable user interface allows clinicians to not only improve the segmentation results but also customize the results based on their judgement and preferences. The results in Tab. 2 and Tab. 3 indicate that the use of the HITL interface is capable of improving the overall performances. Furthermore, AiD Regen only requires users to confirm the suggested boundary on the 2D image, removing the necessity of external training and reducing user’s burdens compared to conventional CAD modeling.

Unlike the conventional 3D scanning devices, AiD Re-

Type	MAE/Std (cm ²)					Accuracy (%)				
	0°	10°	20°	30°	Avr	0°	10°	20°	30°	Avr
A	0.55/0.34	0.37/0.41	0.88/0.87	0.52/0.45	0.59/0.61	95.40	96.91	92.72	95.69	95.15
B	0.19/0.1	0.2/0.08	0.17/0.07	0.26/0.13	0.21/0.11	97.29	97.17	97.66	96.29	97.07
C	2.02/1.53	1.66/1.13	1.51/0.86	2.1/0.77	1.82/1.14	93.15	94.36	94.87	92.89	93.83
Avg	1.2/1.38	0.77/0.97	0.85/0.9	0.96/0.96	0.92/1.05	94.75	96.09	95.08	94.95	95.27

Table 4. The mean absolute error (MAE), standard deviation (Std), and accuracy of the wound surface reconstruction results are shown.

gen uses a single RGB-D image to reconstruct a 3D wound model by taking a hybrid approach that combines 2D segmentation and 3D reconstruction. While such hybrid approach using a single image significantly reduces computational complexities, the results in Tab. 4 indicate that when DFU wounds are captured from 20 cm-distance, AiD Regen adequately reconstructed 3D models with average accuracy of 95.27% and was robust against capturing angles of up to 30°. In the experiment Sec. 4.2, the average surface area of the ground truth wounds was 16.20 cm² and the average MAE was 0.92 cm². Assuming a circular shape, the 95.27% average accuracy indicates the error range from -0.65 mm to 0.63 mm in radius. Since the requirement we empirically received from the clinicians during the clinical trials was ± 1.00 mm in radius, the performance of AiD Regen is acceptable in practice.

AiD Regen’s hybrid approach connects the complex pipeline of multiple well-established techniques including deep learning based semantic segmentation, point cloud processing and filtering, triangular mesh generation, and mesh flattening. This carefully engineered system was able to achieve competitive performances and treat a real patient only with a small and unbalanced 2D image dataset. We hope our system and approach could encourage more researchers to apply established techniques and methods to solve practical real-world problems in various fields.

4.5. Limitations

While AiD Regen’s 3D reconstruction based on a single RGB-D image enhances user experience by simplifying the complex and time-consuming scanning procedures, AiD Regen currently does not support curved wounds that cannot be captured in a single photo. In such cases, users inevitably have to take multiple photos from different angles and generate separate patches per angle. Thus, we are planning to develop a panorama-style capturing method that automatically determines the relationship between multiple images taken from different points-of-view. This way, users will be able to create a patch for complex wound surfaces in a single request.

Though RKT-DFU-800 was collected from various countries without any fixed condition to reflect the real-world environment, it is still not large enough and likely to be biased. That is, our hybrid approach based on the

dataset showed impressive performance but might still produce unexpected results. To handle this limitation, we set up a continuous training that incrementally improves the model as the dataset grows with new wound images. In addition, while clinicians can determine 2D boundaries using our HITL interface, they currently cannot review and modify the final 3D model. Since we want to avoid recreating the burdensome experience of complex CAD-style modeling, we are attempting to develop an augmented reality interface so that users can intuitively manipulate 3D models in the real-world environment.

5. Conclusion

We introduced AiD Regen, a novel system that generates 3D models based on the shape and geometry of wounds combining 2D image segmentation with 3D reconstruction so that they can be printed via 3D bio-printers during the surgery to treat DFUs. AiD Regen seamlessly binds the full pipeline, which includes RGB-D image capturing, semantic segmentation, BGPCP, 3D model reconstruction, and 3D printable G-code generation, into a single system that can be used out of the box. We developed a multi-stage data preprocessing method to handle small and unbalanced DFU image dataset. AiD Regen’s HITL interface enabled clinicians to not only create 3D regenerative patches with a few touch interactions, but also customize and confirm wound boundaries. The results of the experiments indicate that our model outperformed prior wound segmentation models and our reconstruction algorithm was able to generate 3D wound models with compelling accuracy. Our case study on a real DFU patient demonstrated that AiD Regen was effective in treating DFU wounds using a 3D bio-printer.

6. Acknowledgement

We thank Seokhwan You and Soon Yong Kwon for generous support and Gwanmo Park for practical discussion. As participants of this study did not agree for their data to be shared publicly, supporting data is not available.

References

- [1] Maged Aboali, NA Manap, A Majid Darsono, and Zulka-inain Mohd Yusof. Review on three-dimensional (3-d) ac-

- quisition and range imaging techniques. *Int. J. Appl. Eng. Res.*, 12(10):2409–2421, 2017.
- [2] Michael D Abràmoff, Paulo J Magalhães, and Sunanda J Ram. Image processing with imagej. *Biophotonics international*, 11(7):36–42, 2004.
- [3] P Adiewere, Richard B Gillis, S Imran Jiwani, Andrew Meal, Ian Shaw, and Gary G Adams. A systematic review and meta-analysis of patient education in preventing and reducing the incidence or recurrence of adult diabetes foot ulcers (dfu). *Heliyon*, 4(5):e00614, 2018.
- [4] Mario Amrehn, Sven Gaube, Mathias Unberath, Frank Schebesch, Tim Horz, Maddalena Strumia, Stefan Steidl, Markus Kowarschik, and Andreas Maier. Ui-net: Interactive artificial neural networks for iterative image segmentation based on a user model. *VCBM '17*, page 143–147, Goslar, DEU, 2017. Eurographics Association.
- [5] R. Anand, K.G. Mehrotra, C.K. Mohan, and S. Ranka. An improved algorithm for neural network classification of imbalanced training sets. *IEEE Transactions on Neural Networks*, 4(6):962–969, 1993.
- [6] Reza Azad, Maryam Asadi-Aghbolaghi, Mahmood Fathy, and Sergio Escalera. Attention deeplabv3+: Multi-level context attention mechanism for skin lesion segmentation. In *European Conference on Computer Vision*, pages 251–266. Springer, 2020.
- [7] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [8] Yoshua Bengio, Yann Lecun, and Geoffrey Hinton. Deep learning for ai. *Commun. ACM*, 64(7):58–65, June 2021.
- [9] Samuel Budd, Emma C. Robinson, and Bernhard Kainz. A survey on active learning and human-in-the-loop deep learning for medical image analysis. *Medical Image Analysis*, 71:102062, 2021.
- [10] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [11] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L. Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [12] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [13] Ran Cheng, Ryan Razani, Ehsan Taghavi, Enxu Li, and Bingbing Liu. (af)2-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12547–12556, June 2021.
- [14] Yanhua Cheng, Rui Cai, Zhiwei Li, Xin Zhao, and Kaiqi Huang. Locality-sensitive deconvolution networks with gated fusion for rgb-d indoor semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [15] Gary Davidowitz and Philip G Kotick. The use of cad/cam in dentistry. *Dental Clinics*, 55(3):559–570, 2011.
- [16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [17] Mark Everingham, S. M. Ali Eslami, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, Jan 2015.
- [18] Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Distance transforms of sampled functions. *Theory of Computing*, 8(19):415–428, 2012.
- [19] Damir Filko, Robert Cupec, and Emmanuel Karlo Nyarko. Wound measurement by rgb-d camera. *Machine Vision and Applications*, 29(4):633–654, May 2018.
- [20] Manu Goyal, Moi Hoon Yap, Neil D. Reeves, Satyan Rajbhandari, and Jennifer Spragg. Fully convolutional networks for diabetic foot ulcer segmentation. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 618–623, 2017.
- [21] Evelyn Gutierrez, Benjamín Castañeda, Sylvie Treuillet, and Ivan Hernandez. Multimodal and multiview wound monitoring with mobile devices. *Photonics*, 8(10), 2021.
- [22] Xian-Feng Han, Jesse S. Jin, Ming-Jie Wang, and Wei Jiang. Guided 3d point cloud filtering. *Multimedia Tools and Applications*, 77(13):17397–17411, Jul 2018.
- [23] Xian-Feng Han, Hamid Laga, and Mohammed Bennamoun. Image-based 3d object reconstruction: State-of-the-art and trends in the deep learning era. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(5):1578–1604, 2021.
- [24] Daniel A Hashimoto, Guy Rosman, Daniela Rus, and Ozanan R Meireles. Artificial intelligence in surgery: Promises and perils. *Ann. Surg.*, 268(1):70–76, July 2018.
- [25] Caner Hazirbas, Lingni Ma, Csaba Domokos, and Daniel Cremers. Fusetnet: Incorporating depth into semantic segmentation via fusion-based cnn architecture. In *Asian conference on computer vision*, pages 213–228. Springer, 2016.
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1904–1916, 2015.
- [27] Yong He, Hongshan Yu, Xiaoyan Liu, Zhengeng Yang, Wei Sun, Yaonan Wang, Qiang Fu, Yanmei Zou, and Ajmal Mian. Deep learning based 3d segmentation: A survey. *arXiv preprint arXiv:2103.05423*, 2021.
- [28] ROKIT Healthcare. Dr. INVIVO 4D2. <https://rokithealthcare.com/dr-invivo-4d2/>, 2021. Accessed: 2021-11-16.
- [29] ROKIT Healthcare. NewCreatorK. <https://rokithealthcare.com/software-download/>, 2021. Accessed: 2021-11-16.

- [30] Yair Herbst, Shunit Polinsky, Anath Fischer, Yoav Medan, Ronit Schneor, Joshua Kahn, and Alon Wolf. Scan-driven fully-automated pipeline for a personalized, 3D printed low-cost prosthetic hand. In *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*. IEEE, Aug. 2021.
- [31] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [32] Jindong Jiang, Zhijun Zhang, Yongqian Huang, and Lunan Zheng. Incorporating depth into both cnn and crf for indoor semantic segmentation. In *2017 8th IEEE International Conference on Software Engineering and Service Science (ICSESS)*, pages 525–530, 2017.
- [33] Zhizhong Kang, Juntao Yang, Zhou Yang, and Sai Cheng. A review of techniques for 3d reconstruction of indoor environments. *ISPRS International Journal of Geo-Information*, 9(5), 2020.
- [34] Hannah Kerner. Too many ai researchers think real-world problems are not relevant, Sep 2020.
- [35] Rajesh Kesavan, Changam Sheela Sasikumar, V.B. Narayanamurthy, Arvind Rajagopalan, and Jeehee Kim. Management of diabetic foot ulcer with MA–ECM (minimally manipulated autologous extracellular matrix) using 3d bioprinting technology – an innovative approach. *The International Journal of Lower Extremity Wounds*, 0(0):15347346211045625, 0. PMID: 34636693.
- [36] Sung-yeol Kim, Manbae Kim, and Yo-sung Ho. Depth image filter for mixed and noisy pixel removal in rgb-d camera systems. *IEEE Transactions on Consumer Electronics*, 59(3):681–689, 2013.
- [37] Fangzhao Li, Changjian Wang, Xiaohui Liu, Yuxing Peng, and Shiyao Jin. A composite model of wound segmentation based on traditional methods and deep neural networks. *Computational Intelligence and Neuroscience*, 2018:4149103, May 2018.
- [38] Di Lin, Guangyong Chen, Daniel Cohen-Or, Pheng-Ann Heng, and Hui Huang. Cascaded feature network for semantic segmentation of rgb-d images. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [39] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing.
- [40] Hong Liu, Wenshan Wu, Xiangdong Wang, and Yueliang Qian. Rgb-d joint modelling with scene geometric information for indoor semantic segmentation. *Multimedia Tools and Applications*, 77(17):22475–22488, Sep 2018.
- [41] Jing Liu, Yuhang Wang, Yong Li, Jun Fu, Jianguyun Li, and Hanqing Lu. Collaborative deconvolutional neural networks for joint depth estimation and semantic segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 29(11):5655–5666, 2018.
- [42] Kangcheng Liu, Zhi Gao, Feng Lin, and Ben M Chen. Fgnet: Fast large-scale lidar point clouds understanding network leveraging correlated feature mining and geometric-aware modelling. *arXiv preprint arXiv:2012.09439*, 2020.
- [43] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021.
- [44] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [45] PL Muñoz, R Rodríguez, and N Montalvo. Automatic segmentation of diabetic foot ulcer from mask region-based convolutional neural networks. *Journal of Biomedical Research and Clinical Investigation*, 1(1.1006), 2020.
- [46] Nasco. life/form® elderly pressure ulcer foot - light LF00933. <https://www.enasco.com/p/%3Cstrong%3ELife-form%2%AE%3C-strong%3E-Elderly-Pressure-Ulcer-Foot---Light%2BLF00933?searchText=foot%2Bulcer>, 2021. Accessed: 2021-11-16.
- [47] 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. PMID: 31600101.
- [48] Qasem Ramadan and Mohammed Zourob. 3d bioprinting at the frontier of regenerative medicine, pharmaceutical, and food industries. *Frontiers in Medical Technology*, 2:25, 2021.
- [49] Niri Rania, Hassan Douzi, Lucas Yves, and Treuillet Sylvie. Semantic segmentation of diabetic foot ulcer images: Dealing with small dataset in dl approaches. In Abderrahim El Moataz, Driss Mammass, Alamin Mansouri, and Fathallah Nouboud, editors, *Image and Signal Processing*, pages 162–169, Cham, 2020. Springer International Publishing.
- [50] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing.
- [51] Rohan Sawhney and Keenan Crane. Boundary first flattening. *ACM Trans. Graph.*, 37(1), dec 2017.
- [52] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D. Sculley. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. In *NIPS 2017 workshop: Machine Learning for the Developing World*, 2017.
- [53] SHINING 3D. EinScan Pro 2X Handheld Scanner. <https://www.einscan.com/handheld-3d-scanner/einscan-pro-2x/>, 2021. Accessed: 2021-10-29.
- [54] Ke Sun, Yang Zhao, Borui Jiang, Tianheng Cheng, Bin Xiao, Dong Liu, Yadong Mu, Xinggang Wang, Wenyu Liu, and Jingdong Wang. High-resolution representations for labeling pixels and regions. *arXiv preprint arXiv:1904.04514*, 2019.

- [55] Olivier Vincent, Charley Gros, and Julien Cohen-Adad. Impact of individual rater style on deep learning uncertainty in medical imaging segmentation. *arXiv preprint arXiv:2105.02197*, 2021.
- [56] Maximilian Vogt, Adrian Rips, and Claus Emmelmann. Comparison of ipad pro’s lidar and truedepth capabilities with an industrial 3d scanning solution. *Technologies*, 9(2), 2021.
- [57] Kiri Wagstaff. Machine learning that matters. In *ICML*, 2012.
- [58] Chuanbo Wang, D. M. Anisuzzaman, Victor Williamson, Mrinal Kanti Dhar, Behrouz Rostami, Jeffrey Niezgod, Sandeep Gopalakrishnan, and Zeyun Yu. Fully automatic wound segmentation with deep convolutional neural networks. *Scientific Reports*, 10(1):21897, Dec 2020.
- [59] Guotai Wang, Wenqi Li, Maria A. Zuluaga, Rosalind Pratt, Premal A. Patel, Michael Aertsen, Tom Doel, Anna L. David, Jan Deprest, Sébastien Ourselin, and Tom Vercauteren. Interactive medical image segmentation using deep learning with image-specific fine tuning. *IEEE Transactions on Medical Imaging*, 37(7):1562–1573, 2018.
- [60] Roman V Yampolskiy and MS Spellchecker. Artificial intelligence safety and cybersecurity: A timeline of ai failures. *arXiv preprint arXiv:1610.07997*, 2016.
- [61] Leila Yazdanpanah, Morteza Nasiri, and Sara Adarvishi. Literature review on the management of diabetic foot ulcer. *World journal of diabetes*, 6(1):37, 2015.
- [62] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip H.S. Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16259–16268, October 2021.
- [63] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [64] Y. M. Zhao, Edward H. Currie, Louis Kavoussi, and Sina Y. Rabbany. Laser scanner for 3d reconstruction of a wound’s edge and topology. *International Journal of Computer Assisted Radiology and Surgery*, Aug 2021.
- [65] Michael Zollhöfer, Patrick Stotko, Andreas Görlitz, Christian Theobalt, Matthias Nießner, Reinhard Klein, and Andreas Kolb. State of the art on 3d reconstruction with rgb-d cameras. *Computer Graphics Forum*, 37:625–652, 05 2018.