

WoundAmbit: Bridging State-of-the-Art Semantic Segmentation and Real-World Wound Care

Vanessa Borst¹ (✉), Timo Dittus¹, Tassilo Dege²,
Astrid Schmieder², and Samuel Kounev¹

¹ University of Würzburg, 97070 Würzburg, Germany
{vanessa.borst,timo.dittus,samuel.kounev}@uni-wuerzburg.de

² University Hospital of Würzburg, 97070 Würzburg, Germany
{Dege_T, Schmieder_A}@ukw.de

Abstract. Chronic wounds affect a large population, particularly the elderly and diabetic patients, who often exhibit limited mobility and co-existing health conditions. Automated wound monitoring via mobile image capture can reduce in-person physician visits by enabling remote tracking of wound size. Semantic segmentation is key to this process, yet wound segmentation remains underrepresented in medical imaging research. To address this, we benchmark state-of-the-art deep learning models from general-purpose vision, medical imaging, and top methods from public wound challenges. For a fair comparison, we standardize training, data augmentation, and evaluation, conducting cross-validation to minimize partitioning bias. We also assess real-world deployment aspects, including generalization to an out-of-distribution wound dataset, computational efficiency, and interpretability. Additionally, we propose a reference object-based approach to convert AI-generated masks into clinically relevant wound size estimates and evaluate this, along with mask quality, for the five best architectures based on physician assessments. Overall, the transformer-based TransNeXt showed the highest levels of generalizability. Despite variations in inference times, all models processed at least one image per second on the CPU, which is deemed adequate for the intended application. Interpretability analysis typically revealed prominent activations in wound regions, emphasizing focus on clinically relevant features. Expert evaluation showed high mask approval for all analyzed models, with VWFormer and ConvNeXtS backbone performing the best. Size retrieval accuracy was similar across models, and predictions closely matched expert annotations. Finally, we demonstrate how our AI-driven wound size estimation framework, *WoundAmbit*, is integrated into a custom telehealth system. Our code and supplementary material are available on GitHub and Zenodo, respectively.

Keywords: Semantic Segmentation · Benchmarking · Deep Learning · CNN · Transformer · Clinical Application · Tele-Medicine · Wound Care.

1 Introduction

Wound care is a critical aspect of healthcare, particularly for chronic wounds that require ongoing monitoring and treatment. Current clinical practice often

relies on manual wound measurements, such as estimating wound area by multiplying its longest length by its largest perpendicular width using a ruler or metric tape—although this often leads to overestimation due to irregular wound shapes [26]. Langemo et al. further reported differing interpretations among clinicians regarding how to define and measure wound length and width [16]. This lack of standardization contributes to measurement subjectivity and may limit comparability between assessments. An alternative approach involves tracing the wound on a transparent film and estimating the area using a metric grid, which may offer improved measurement reliability but still suffers from subjectivity in boundary delineation and partial cell interpretation [26,16]. Beyond potential inconsistencies, manual assessments are often invasive, relying on either proximity to or direct physical contact with the wound. This may cause patient discomfort and requires in-person visits with healthcare professionals (HCPs), posing logistical challenges for both patients and providers. Automated wound segmentation offers a promising alternative, as AI-driven wound size estimation from RGB images can be integrated into various systems, particularly mobile phones with cameras. This enhances accessibility, enabling remote wound monitoring from home without requiring specialized hardware. Unlike tracing methods, AI-based approaches are non-invasive, eliminating direct wound contact.

This shift toward automation aligns with broader advancements in AI-driven semantic segmentation, which have mainly been shaped by the evolution of convolutional neural networks (CNNs) and the emergence of vision transformers (ViTs). Thisanke et al. [29] review modern transformer-based approaches, whereas Minaee et al. [20] focus on deep learning (DL) for image segmentation excluding ViTs. Recently, (multi-modal) vision foundation models, such as Segment Anything [15], have gained attention due to their large-scale pretraining and versatility across different downstream tasks [1]. Medical image segmentation has also transitioned towards DL. Azad et al. present a comprehensive review of recent advances using ViTs [2], while Rayed et al. [23] review DL approaches more broadly. Additionally, foundation models like Segment Anything have been adapted for medical imaging, as demonstrated by MedSAM [19].

In contrast, wound segmentation remains relatively underexplored, partly due to the scarcity of relevant datasets. Notable exceptions with more than 1,000 wound images include FUSeg [31] and DFUC [14], both of which provide annotated images of diabetic foot ulcer (DFU) for public challenges. Additionally, Oota et al. [21] provide a dataset with 2,686 images of eight wound types, where annotations extend beyond the wound to include peri-wound skin areas. Regarding existing wound segmentation techniques, early methods focused primarily on traditional feature engineering-based machine learning [32]. Over time, DL approaches such as WNet [21], FUSegNet [6] and other CNN-based techniques [30,17,18,5,8] have emerged, alongside a few approaches for interactive wound segmentation [37]. However, certain limitations remain:

1. Limited Adoption of SotA Vision Models for Wound Analysis:

Despite advancements in computer vision, investigations into the suitability of state-of-the-art (SotA) models, particularly ViTs, for wound segmentation re-

main limited. A 2022 survey on DL for wound analysis [38] did not mention any transformer-based segmentation methods, and recent breakthroughs in general-purpose (GP) vision models have yet to be applied to wound analysis.

2. Lack of Consideration for Practical Deployment: Few studies evaluate efficiency metrics such as GMACs/FLOPs alongside segmentation performance [30,21], and few report inference times or latency [37,21]. However, computational complexity is a critical factor for resource-constrained settings such as the medical sector, where GPU availability is limited. With few exceptions, such as the visualization of feature maps from different layers [37], explainable AI techniques are rarely applied despite their potential to improve trust among clinicians. Lastly, model generalizability to out-of-distribution (OOD) data is often overlooked, even though assessing robustness is crucial due to variability in wound types, skin tones, and uncontrolled home-based monitoring conditions such as lighting, background, and hardware.

3. Gap Between Segmentation and Clinically Relevant Wound Size: Existing public challenges (FUSeg, DFUC) and most wound segmentation approaches [6,30,17,18,8] do not address the conversion of segmentation masks into real-world wound size. Exceptions include Wang et al. [32], who rely on a specialized imaging box, and Chairat et al. [4], who use a custom calibration chart with a U-Net-based model incorporating EfficientNet/MobileNetv2 encoders for wound segmentation—but do not assess size retrieval accuracy. Chino et al. [5] incorporate measurement ruler and tape detection as a third class in a U-Net-based model and combine it with pixel density estimation to determine wound area. Similarly, Foltynski and Ladyzynski [7] train a CNN to detect both wounds and dual calibration markers that need to be placed below and above the wound. Proprietary solutions like Swift Skin & Wound [22] and imitoWound [12] exist but provide little insight into their calibration methods and algorithms.

To address these limitations, we introduce *WoundAmbit*, an end-to-end solution for automated wound size estimation from RGB images that bridges the gap between modern DL and practical wound care. To the best of our knowledge, this study is the first to systematically transfer a diverse set of SotA GP vision models to the wound domain and benchmark them against both medical and wound-specific segmentation methods within a unified evaluation framework. As summarized in Figure 1, our key contributions are two-fold:

1. Comprehensive DL Benchmark: We conduct a rigorous benchmarking study by systematically selecting 12 SotA DL architectures from wound-specific, medical, and GP vision models, covering a diverse range of design paradigms, including CNNs, ViTs, and hybrid models. To ensure comparability, all models are trained under standardized conditions on publicly available data and evaluated using 5-fold cross-validation (CV). In addition, our evaluation framework emphasizes clinically relevant properties by assessing generalizability on a dedicated OOD dataset. Specifically created for this work, the dataset comprises 343 wound images taken at various body sites. Moreover, we analyze computational efficiency, including trainable parameters, Giga Multiply-Accumulate Operations (GMACs), and inference times on both GPU and CPU. Lastly, we inves-

tigate model interpretability using Gradient-weighted Class Activation Mapping (Grad-CAM) visualizations to assess clinically relevant decision-making.

2. Real-World Deployment: To translate AI-generated wound masks into clinically meaningful size estimates, we develop and validate a reference object (RO)-based approach for precise wound surface measurement. Unlike existing methods that require neural networks to segment calibration stickers [5,7], our approach leverages ArUco marker detection, making it independent of the AI algorithm used for segmentation. To evaluate reliability, we construct a dataset of 20 diverse wound images and obtain expert assessments of AI-generated mask quality from three dermatologists. Additionally, we compare AI-derived wound size estimates with physician annotations to quantify measurement accuracy. Finally, we propose a practical integration strategy for embedding AI-driven wound size estimation into a custom telemedicine framework for remote monitoring.

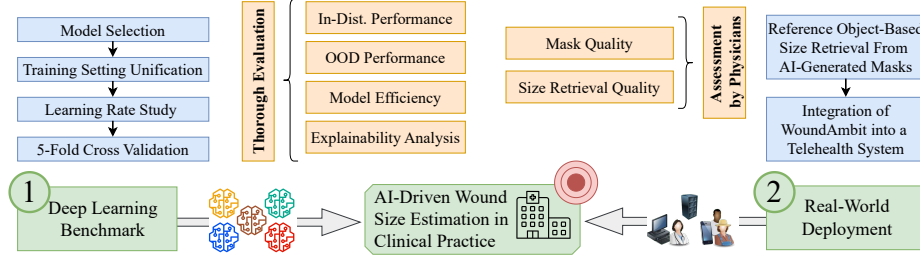


Fig. 1: Schematic visualization of our *WoundAmbit* approach.

The paper is structured as follows: Section 2 describes the DL benchmark (Fig. 1, left), covering model selection, methodology, datasets, and results. Section 3 details the size retrieval process and its evaluation, along with the integration of *WoundAmbit* into a specially designed telehealth system (Fig. 1, right). Finally, Section 4 discusses key findings, and Section 5 concludes the study.

2 Deep Learning Benchmark

2.1 Model Selection

Given the large number of novel methods introduced in recent years, it is impractical to include all current SotA models in this study. To ensure a benchmarking process that is as representative and balanced as possible, we devise four categories, selecting several representatives from each of them as follows:

(I) *DFU segmentation*: To specifically address the challenges of wound segmentation (WS), we include the best-performing models from two publicly available DFU segmentation challenges, as they are directly tailored to the domain of interest (HarDNet-DFUS [17], FUSegNet [6]). (II) *Medical segmentation*: Medical segmentation (MS) architectures, such as U-Net and its variants, have demonstrated broad applicability across diverse medical imaging tasks [28]. To assess

the potential of recent architectures that have shown effectiveness beyond wound care, we select three models with strong performance in medical image segmentation outside the wound domain (FCBFormer [25], HiFormer-B [10], MISSFormer [11]). (III) *GP SS models*: To incorporate cutting-edge developments, we include two widely recognized GP semantic segmentation (SS) models (SegFormer-B3 [35] and SegNeXt-L [9]), along with a recent multi-scale decoder (VWFormer [36]) for SS. The VWFormer model is integrated with two different backbones, MiT-B3 and ConvNeXt-S, denoted as VW-MiT and VW-Conv, respectively. (IV) *GP vision models*: To account for advanced (multi-task) vision models (VM), we select two promising approaches from this category (InternImage-T [33], TransNeXt-Tiny [27]). As common in literature, we adapt them for segmentation by integrating a UPerHead [34] decoder. (V) *Baseline*: Apart from these categories, we include U-Net [24], a proven and widely used model in biomedical segmentation, as a baseline.

Table 1 summarizes our final selection, striking a balance between domain specificity, cross-domain variety, and methodological diversity. The key ideas of all architectures are briefly summarized in Section A.1, with further details available in the corresponding publications. Overall, our selection process is guided by the following criteria: (I) *Architectural diversity*: To ensure a comprehensive representation of various design paradigms, we aim for a balance between CNN-based, ViT-based, and hybrid methods. (II) *Computational efficiency*: Given the limited resources in healthcare settings, we prioritize models that offer a good trade-off between performance and computational feasibility. For GP models, we select architectures with approximately 50–60M parameters, while wound-specific and medical models range from 30–70M parameters, mainly due to the limited availability of varied model sizes in these domains. (III) *Scientific impact and recognition*: We select models with significant visibility in their respective fields, prioritizing those published in high-impact conferences and journals. (IV) *Code availability*: To minimize the risk of implementation errors, we limit our selection to models with publicly available code.

Table 1: Overview of the selected methods besides our baseline U-Net [24].

Healthcare-Related					General-Purpose				
Cat.	Model	Type	Size	Y.	Cat.	Model	Type	Size	Y.
MS	FCBFormer [25]	Hybrid	53M	'22	SS	SegFormer-B3 [35]	Transf.	48M	'21
	HiFormer-B [10]	Hybrid	30M	'23		SegNeXt-L [9]	CNN	49M	'22
	MISSFormer [11]	Transf.	43M	'23		2× VWFormer [36]	-Depends ¹	'24	
WS	HarDNet-DFUS [17]	CNN	52M	'22	VM	InternImage-T [33] ²	CNN	58M	'23
	FUSegNet [6]	CNN	71M	'24		TransNeXt-Tiny [27] ²	Transf.	58M	'24

¹ 51M with MiT-B3/ 57M with ConvNext-S

² Parameter count includes UPerHead decoder

2.2 Unified Training Procedure: Methodological Details

We establish a unified benchmarking environment by standardizing the training process as follows: To ensure an unbiased comparison while enabling the use of pre-trained weights, the deep learning methods were implemented with their architecture-specific settings, as recommended in their respective official publications and code repositories. Specifically, we use ImageNet-pretrained backbones for all methods. While maintaining key architectural parameters (e.g., layer configurations and activation functions), we standardize other training settings across all models. These include input tensor dimensions, the number of training epochs, early stopping criteria, optimizer type, loss function, and the data augmentation pipeline. Further details on preprocessing, data augmentation, and exact training configurations are provided in Section A.2. To minimize biases associated with fixed learning rates, we conduct a preliminary hyperparameter tuning step for each architecture, optimizing the learning rate within a predefined search range (10^{-4} , 5×10^{-5} , 10^{-5}). The best-performing learning rate for each method (see Section A.2) is then used for subsequent five-fold CV, while all other training settings remain unified. Overall, this ensures that performance differences arise from the intrinsic capabilities of the architectures rather than variations in training configurations, particularly data augmentation.

2.3 Standardized Evaluation Procedure

Our assessment strategy employs a unified evaluation framework, incorporating both traditional segmentation metrics and practical aspects essential for clinical applications. For *segmentation performance*, we use metrics such as mean Intersection over Union (mIoU), Dice Similarity Coefficient (mDSC), precision (mPrc), and recall (mRec). These metrics evaluate performance on both our main dataset and unseen OOD data, with the latter specifically assessing the *generalization capability* across diverse wound types. In terms of *model efficiency*, we report the number of trainable parameters and GMACs, along with mean inference time for GPU and CPU execution and throughput in images per second (IPS). Finally, we assess *explainability* using *Grad-CAM*-based visualizations. Further information and implementation details are available in Section A.3.

2.4 Datasets

Our experiments mainly rely on two datasets, detailed in Section A.4. The first, *CFU*, is used for model training, including learning rate studies and final CV. It consists of a custom combination of the publicly available DFUC’22 dataset [14], which contains 2,000 annotated images, and the FUSeg’21 challenge dataset [31], with 1,010 labeled images. After removing duplicates and highly similar images, 2,887 unique images remain for our experiments. Additionally, we use the DFUC’22 test set for external validation of our models via the challenge’s live leaderboard. The second dataset, denoted as *out-of-distribution (OOD)*, was collected at the University Hospital of Würzburg with approval from the local

ethics committee. It comprises 343 expert-annotated wound images from various anatomical sites, extending beyond foot ulcers. To better reflect real-world clinical conditions, where patient-acquired images often lack standardized imaging protocols, certain wounds were intentionally captured multiple times from varying distances, angles, and perspectives. Notably, *OOD* is used exclusively for evaluation, not for training. Due to privacy regulations, it remains confidential; however, selected examples are shared at GitHub with written consent.

2.5 Performance on CFU Dataset and DFUC’22 Live Leaderboard

Table 2 presents the 5-fold CV results on CFU alongside each model’s performance on the DFUC’22 live leaderboard. For the latter, instead of selecting the best checkpoint from individual folds, segmentation masks are generated using pixel-wise majority votes (PMVs) across all five instances of each architecture. This approach mitigates overfitting to individual training folds while leveraging the collective strengths of multiple trained instances.

On CFU, TransNeXt demonstrated the highest performance, achieving a mIoU of 79.8 and an mDSC of 88.7. SegNeXt followed closely with a mIoU of 79.5 and an mDSC of 88.6, exhibiting strong consistency across the CV folds. With similarly low inter-fold variability, SegFormer ranked third, achieving a mIoU of 78.9 and an mDSC of 88.2. On the DFUC’22 leaderboard, TransNeXt also achieved the highest performance among our models, with a mIoU of 62.8 and an mDSC of 73.0, slightly surpassing SegNeXt and VW-Conv—the latter ranking third despite its moderate performance on CFU. Notably, TransNeXt ranked 6th out of 60 (top 10%) based on the best submission per participant despite no dataset-specific optimization, which highlights its strong out-of-the-box performance. In contrast, U-Net and most medical approaches, except FCBFormer, showed lower segmentation performance and ranked further down on both CFU and DFUC’22.

Table 2: 5-fold CV (mean±SD, %) and ensemble DFUC’22 leaderboard scores.

Type	Avg. CFU				DFUC’22	
	mIoU↓	mDSC	mPrc	mRec	mIoU	mDSC
TransNeXt	79.8±1.4	88.7±0.8	90.9±0.4	86.7±1.8	62.8	73.0
SegNeXt	79.5±0.7	88.6±0.4	90.7±0.6	86.6±0.6	62.1	72.3
SegFormer	78.9±0.8	88.2±0.5	90.4±0.7	86.1±1.4	61.8	72.1
FCBFormer	78.6±1.5	88.0±0.9	90.6±0.6	85.6±1.9	62.0	72.2
InternImage	78.5±0.9	88.0±0.5	89.8±1.5	86.2±0.6	61.7	72.0
VW-MiT	78.5±1.7	87.9±1.0	90.2±1.1	85.8±2.2	61.7	72.0
VW-Conv	78.4±0.5	87.9±0.3	90.1±1.9	85.9±2.0	62.0	72.3
FUSegNet	78.0±1.3	87.6±0.8	90.6±0.7	84.9±1.3	61.3	71.6
HarDNet	76.9±1.4	86.9±0.9	88.3±2.4	85.8±2.5	60.4	70.8
U-Net	74.1±0.9	85.1±0.6	88.1±1.1	82.5±1.5	57.6	68.0
MISSFormer	70.0±2.7	82.3±1.8	85.8±3.4	79.3±3.3	55.6	66.4
HiFormer	73.8±1.8	84.9±1.2	87.6±1.6	82.5±1.7	57.7	68.2

2.6 Performance on Out-of-Distribution (OOD) Data

Table 3 presents the segmentation performance on the unseen OOD dataset, which includes previously unobserved anatomical regions, such as the head and breast, thereby assessing model generalization to domain shifts. In addition to reporting the average and standard deviation (SD) across the five CV models per architecture, we again provide PMV results to ensure a more stable and unbiased evaluation. Notably, for all models, the PMV performance surpasses the average performance of the best models from individual folds in both mIoU and mDSC, further confirming its robustness. TransNeXt demonstrated the strongest generalization, achieving the highest mIoU (79.4) and mDSC (88.5) in the PMV setting, followed closely by InternImage (mIoU 78.8, mDSC 88.1) and VW-MiT (mIoU 78.0, mDSC 87.6). The performance drop relative to the in-distribution CFU data was minimal ($\leq 0.5\text{pp}$ mIoU, $\leq 0.3\text{pp}$ mDSC) for the top three models. InternImage even showed a 0.3pp improvement in mIoU, indicating strong adaptability to novel wound types. SegFormer, VW-Conv, and HardNet also maintained competitive performance relative to CFU ($\sim 1\text{pp}$ mIoU drop). FCBFormer achieved the highest absolute PMV scores among the medical architectures (mIoU 76.5, mDSC 86.7). In contrast, the remaining medical models, except for HardNet, exhibited substantial performance declines (4–7.9pp mIoU), with HiFormer dropping below 68% mIoU. Notably, SegNeXt, despite its strong in-distribution and DFUC’22 performance, saw a sharp decline of 3.7pp in mIoU.

Table 3: Segmentation performance (mean \pm SD, %) on the OOD dataset using the five trained CV models, with CFU mean IoU and DSC for reference.

Type	Avg. CFU		Avg. OOD				Maj. vote OOD			
	mIoU	mDSC	mIoU	mDSC	mPrec	mRec	mIoU \downarrow	mDSC	mPrec	mRec
TransNeXt	79.8	88.7	78.3\pm1.7	87.8\pm1.1	93.6 \pm 0.2	82.7\pm1.9	79.4	88.5	94.3	83.4
InternImage	78.5	88.0	76.0 \pm 3.2	86.3 \pm 2.1	92.4 \pm 0.6	81.1 \pm 3.6	78.8	88.1	93.7	83.2
VW-MiT	78.5	87.9	75.7 \pm 1.6	86.2 \pm 1.0	92.8 \pm 0.5	80.5 \pm 2.0	78.0	87.6	93.7	82.3
SegFormer	78.9	88.2	76.3 \pm 1.2	86.5 \pm 0.7	93.1 \pm 0.6	80.8 \pm 1.4	77.7	87.5	94.1	81.7
VW-Conv	78.4	87.9	74.9 \pm 2.7	85.7 \pm 1.8	93.1 \pm 0.8	79.4 \pm 3.5	77.2	87.1	94.4	80.9
FCBFormer	78.6	88.0	74.0 \pm 1.9	85.1 \pm 1.3	93.2 \pm 1.0	78.3 \pm 2.7	76.5	86.7	94.5	80.0
HardNet	76.9	86.9	73.0 \pm 2.6	84.4 \pm 1.7	91.1 \pm 2.1	78.7 \pm 4.1	75.9	86.3	93.1	80.4
SegNeXt	79.5	88.6	74.5 \pm 1.7	85.4 \pm 1.1	93.4 \pm 0.3	78.6 \pm 2.0	75.8	86.2	94.4	79.4
FUSegNet	78.0	87.6	71.5 \pm 3.2	83.4 \pm 2.2	94.2\pm0.8	74.9 \pm 3.9	74.0	85.1	95.7	76.5
U-Net	74.1	85.1	65.6 \pm 2.7	79.2 \pm 2.0	91.9 \pm 0.6	69.7 \pm 3.3	67.5	80.6	93.3	70.9
MISSFormer	70.0	82.3	64.3 \pm 2.6	78.2 \pm 2.0	88.7 \pm 1.6	70.1 \pm 3.5	66.0	79.5	90.2	71.1
HiFormer	73.8	84.9	63.1 \pm 5.1	77.3 \pm 3.9	90.8 \pm 2.0	67.7 \pm 6.3	65.9	79.4	92.4	69.6

2.7 Model Efficiency Analysis

We report key computational metrics in Table 4. Among all architectures, FUSegNet has the highest number of trainable parameters (71M), while HiFormer and

U-Net are the smallest. In terms of GMACs, MISSFormer is the most efficient, whereas U-Net, and, especially, TransNeXt and InternImage exhibit higher computational complexity. Interestingly, parameter count and GMACs do not always directly correspond to inference time and throughput (TP). For instance, despite having similar parameter counts to most models (50–60M), TransNeXt and FCBFormer exhibit considerably higher inference times and lower TP, making them the slowest on both GPU and CPU. In contrast, U-Net, despite having the third-highest GMAC count, achieves the fastest GPU inference time, highlighting the influence of architectural design and optimizations on computational efficiency. The results reveal a trade-off between efficiency and segmentation performance. While models such as U-Net, MISSFormer, and HiFormer excel in inference speed on both GPU and CPU, advanced vision models like TransNeXt and InternImage offer superior segmentation performance at a moderate computational cost. Notably, even TransNeXt, the slowest model, maintains a throughput of approximately one IPS on the CPU and up to 24 IPS on the GPU.

Table 4: Evaluation of model efficiency, with OOD maj. vote mIoU for reference.

Type	OOD	Params	GMACs	\varnothing Inference Time \pm SD (in ms)		TP (in IPS)	
	mIoU \downarrow	(in M)		GPU	CPU	GPU	CPU
TransNeXt	79.4	57.74	238.41	41.58 \pm 0.10	1012.75 \pm 243.11	24.05	0.99
InternImage	78.8	58.37	234.23	28.71 \pm 0.29	692.35 \pm 42.90	34.83	1.44
VW-MiT	78.0	51.42	62.79	25.65 \pm 0.11	439.98 \pm 55.75	38.98	2.27
SegFormer	77.7	47.22	71.18	23.54 \pm 0.08	429.89 \pm 17.95	42.48	2.33
VW-Conv	77.2	57.00	77.43	25.55 \pm 0.03	339.08 \pm 38.59	39.13	2.95
FCBFormer	76.5	52.96	149.93	51.27 \pm 0.09	875.13 \pm 41.01	19.51	1.14
HarDNet	75.9	51.06	138.92	38.01 \pm 0.16	669.29 \pm 17.98	26.31	1.49
SegNeXt	75.8	48.77	64.18	29.50 \pm 0.42	616.49 \pm 59.53	33.90	1.62
FUSegNet	74.0	70.97	35.51	32.99 \pm 0.13	526.02 \pm 69.17	30.31	1.90
U-Net	67.5	31.03	192.67	13.79 \pm 0.02	307.96 \pm 15.82	72.50	3.25
MISSFormer	66.0	42.46	7.16	17.78 \pm 0.15	203.07 \pm 19.05	56.24	4.92
HiFormer	65.9	25.51	11.51	14.09 \pm 0.13	251.57 \pm 19.80	70.96	3.98

2.8 Explainability Insights

To assess the decision-making of our models, we applied *Grad-CAM* to generate visual explanations of predictions. Figure 2 presents visualizations for the top three and two lowest-performing models from Table 3, along with the predicted segmentation masks (red) for selected OOD images. Additional comparisons across architectures and further examples of different body localization, wound size, and color variations are provided in Section A.5.

By comparing heatmaps with ground truth (GT) masks (green), we qualitatively assess model reliability, focusing on whether activations align with clinically relevant wound features rather than artifacts. Across all models, strong activations were frequently observed in wound regions, aligning well with GT masks.

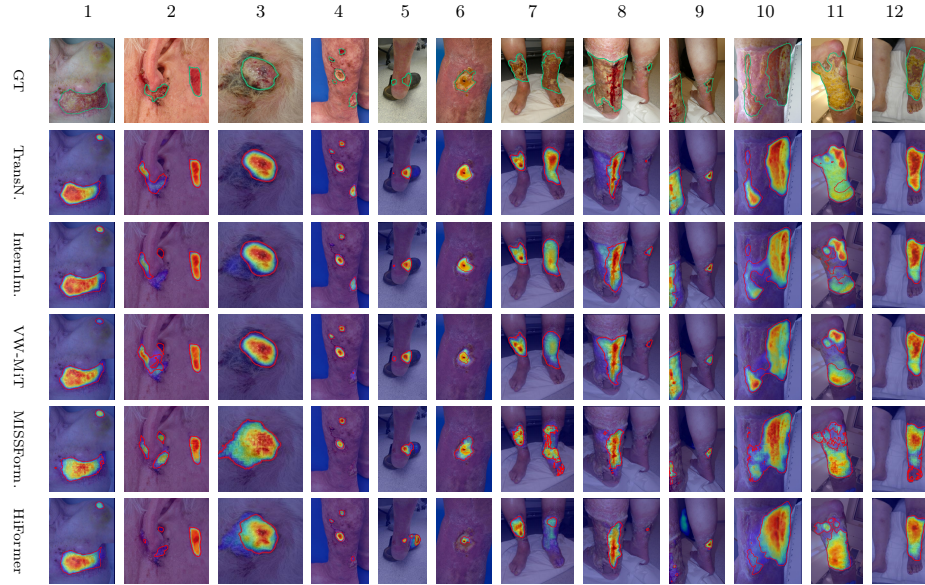


Fig. 2: *Grad-CAM* for exemplary OOD images, alongside their GT annotation.

Additionally, non-wound regions often exhibited minimal to no activation, suggesting effective prioritization of pathological features. Qualitative comparisons indicate that higher-performing models from Table 3 tend to generate sharper heatmap boundaries and achieve more precise wound localization, whereas lower-performing models exhibit weaker or more diffuse activations (3, 7). In some cases, non-wound objects (e.g., shoes, 5) were misidentified as wounds, while certain wound regions (6, 7, 8), including those near image edges (9), were not fully captured, especially by lower-performing models. Additionally, weaker models occasionally generated pixelated or irregular mask predictions, as indicated by frayed or poorly defined red contours (7, 11, 12). While not necessarily generalizable, these qualitative insights align with quantitative performance metrics and enhance interpretability for clinicians, allowing for a deeper understanding of model behavior and fostering greater trust in AI systems.

3 Real-World Deployment

3.1 Automated Size Retrieval from AI-Generated Masks

Reference Object (RO) Design. Based on Chairat et al. [4], we developed a custom RO with slight adjustments to size and layout (see Fig. 3a). Our RO is placed next to the wound and features four ArUco markers, each with a 12 mm side length and a 4 mm margin of white space to facilitate accurate detection. To enhance functionality, we included a Macbeth color chart with 24 color patches

between the markers, aligned such that the top and bottom edges of the patch groups align with the marker corners to simplify automated extraction.

Benefits. Unlike existing methods [5,7], our ArUco-based RO, with its known dimensions, enables precise object size estimation independent of AI algorithms. Notably, this approach eliminates the need for including RO detection as an additional class in model training. For future applications, the Macbeth color chart enables color calibration and wound color assessment. In environments with few natural reference points (e.g., human skin), the fixed design helps estimate camera position and rotation, which could be used for real-time patient guidance during image capture (e.g., adjusting distance or angle).

Size Retrieval Process. We detect the ArUco markers using OpenCV’s `ArucoDetector`. To estimate the pixel-to-millimeter (px/mm) ratio, we design a robust approach that utilizes multiple ArUco markers whenever possible, computing the ratio using predefined real-world distances between specific marker pairs. By calculating the mean px/mm ratio across all available pairs, we reduce the impact of perspective distortion and measurement noise. Moreover, this approach enables more stable estimation, even in cases where some markers are partially occluded or missing. If only one marker is detected, we estimate the px/mm ratio by averaging its width and height. Afterward, we convert segmented wound areas into square millimeters using this ratio. In addition, we determine the wound’s height and width following standard clinical measurement practices [16]. We identify the longest diagonal of each mask contour (with at least seven points). Then, a perpendicular vector is computed and moved incrementally along the diagonal. At each step, we measure the distance between intersections of the perpendicular line and the contour, selecting the two points with the greatest separation to define the second diagonal, representing wound width. Figure 3 demonstrates the size retrieval process for three representative patients. Green contours indicate detected markers and wound regions, while pink lines represent the diagonals approximating wound height and width. Red numbers show the retrieved marker IDs. Notably, accurate wound size estimation requires the RO to be placed as close as possible to the wound and, crucially, within the same plane to minimize distortion. In contrast, Subfigure 3c shows a failure case where the wound is identified correctly, but its size is underestimated due to the markers being positioned closer to the camera than the wound.




Fig. 3: Visualization of the size retrieval process for exemplar real-world patients.

Evaluation. To evaluate performance, we collect $N = 20$ diverse wound images using our RO and an iPhone 11. The top five methods, selected based on PMV mIoU scores from the OOD dataset, are assessed on these images using PMVs per architecture. Additionally, an ensemble of all $5 \times 5 = 25$ models is constructed. For qualitative evaluation, $|D| = 3$ dermatologists independently review AI-generated masks. A total of 120 image-mask pairs ($N \times 6$ AI variants) are presented in random order, with each mask rated as either “good” or “bad”. Physicians also estimate wound sizes and select the best mask per image. Figure B.1 provides screenshots of the annotation tool. We define the following for evaluation: The *Clinical Mask Approval (CMA)* measures how often physicians rate the masks as “good” across all images and evaluators. The *Expert Choice Rate (ECR)* quantifies the proportion, averaged across all images and dermatologists, where a model’s mask is selected as the best of six for a given image.

$$\text{CMA} = \frac{\# \text{ Good}}{|D| \times N} \times 100 \quad (1) \quad \text{ECR} = \frac{\# \text{ Times selected as best}}{|D| \times N} \times 100\% \quad (2)$$

Since physicians provide separate height and width estimates instead of direct area measurements, we use these as proxies for size estimation accuracy. The GT height and width for image i are defined as the **mean** of all $|D|$ physician estimates, denoted as $H_{GT}^{(i)}$ and $W_{GT}^{(i)}$. To assess size retrieval quality, we employ the Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE). For each dimension $X \in \{H, W\}$, these metrics quantify the deviation between model predictions $X_{pred}^{(i)}$ and the expert-derived GT $X_{GT}^{(i)}$:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |X_{pred}^{(i)} - X_{GT}^{(i)}| \quad (3) \quad \text{MAPE} = \frac{1}{N} \sum_{i=1}^N \frac{|X_{pred}^{(i)} - X_{GT}^{(i)}|}{X_{GT}^{(i)}} \times 100\% \quad (4)$$

However, size estimation in medical imaging is inherently subjective, leading to significant variability among physicians [13]. To ensure a meaningful evaluation, we quantify inter-rater variability using a relative deviation metric. We exclude any image i where the following deviation exceeds a threshold of 0.5 in either height or width annotations, with $X \in H, W$ representing the dimension, and $X_{d_j}^{(i)}$ denoting the size annotation provided by physician d_j for image i :

$$\text{Relative Deviation} = \frac{\max(X_{d_1}^{(i)}, X_{d_2}^{(i)}, X_{d_3}^{(i)}) - \min(X_{d_1}^{(i)}, X_{d_2}^{(i)}, X_{d_3}^{(i)})}{\text{median}(X_{d_1}^{(i)}, X_{d_2}^{(i)}, X_{d_3}^{(i)})} \quad (5)$$

After removing inconsistently annotated images, we use the remaining seven for error calculations (see Table 5). Figure 4 confirms their variability in wound shape, size, and skin tone, ensuring a representative subset for size retrieval evaluation. We also report the mean predicted height and width (MPH, MPW) with SD across the $N = 7$ images, alongside the average wound dimensions estimated by the $|D|$ physicians. This analyzes model consistency across wound types and retrieval capabilities relative to expert estimates. For completeness, Table B.1 includes the masks, raw size predictions (including area), and expert estimates for all 20 images, which are also shared at GitHub with written consent.



Fig. 4: Overview of the seven wounds with consistent expert size estimates.

The CMA ranged from 83.3% to 86.7% across models, reflecting high mask quality and strong physician approval. ECR varied notably, with VW-Conv achieving the highest approval rate (35.0%) and TransNeXt and the Ensemble the lowest (8.3%). For size retrieval, all models performed similarly, with MPH and MPW closely aligning with the GT. On average, height estimates were slightly overestimated, while width was slightly underestimated. Nevertheless, consistently low MAE (≤ 4.4 mm height, ≤ 3.2 mm width) and MAPE ($\leq 14.1\%$ height, $\leq 16.2\%$ width) indicate robust size retrieval across wound sizes.

Table 5: Expert-based mask and size retrieval assessment.

Model	Mask Quality		Size Retrieval Quality ($N = 7$)					
	$(N = 20)$		Height			Width		
	CMA ¹	ECR ¹	MPH ^{2,3}	MAE ²	MAPE ¹	MPW ^{2,4}	MAE ²	MAPE ¹
TransNeXt	85.0	8.3	54.6 \pm 41.0	4.4	12.3	27.2 \pm 15.6	3.2	15.8
InternImage	86.7	15.0	54.9 \pm 40.9	3.9	12.2	26.6 \pm 15.0	3.2	16.2
VW-MiT	86.7	13.3	55.4 \pm 41.5	4.3	14.1	27.7 \pm 16.0	2.3	13.3
SegFormer	83.3	13.3	54.8 \pm 41.2	3.8	12.0	27.1 \pm 15.8	2.6	13.2
VW-Conv	85.0	35.0	54.8 \pm 41.6	4.0	9.9	26.6 \pm 15.8	3.2	14.4
Ensemble	86.7	8.3	55.0 \pm 41.3	4.1	12.2	27.0 \pm 15.7	2.9	15.1

¹ In % ² In mm

³ Mean H_{GT} : 54.3 \pm 41.6

⁴ Mean W_{GT} : 29.2 \pm 16.3

3.2 Seamless Integration into a Custom Telehealth System

We integrated the AI module, which demonstrated superior performance in terms of ECR, VW-Conv, and wound size retrieval, into a custom-developed telehealth platform to enable future evaluation of its impact on patient outcomes. The system architecture (see Figure 5) consists of three core components: (1) the *WoundAIssist* mobile application [3], (2) a web interface for physicians, and (3) a backend with a dedicated microservice for AI processing. Among other features, patients use our purpose-built app, *WoundAIssist*, to capture and upload wound images at regular intervals, accompanied by self-reported data on wound-specific factors (e.g., pain, itching, oozing) and overall well-being (e.g., mood, activity impact, quality of life) [3]. HCP use a web-based dashboard for remote patient monitoring, providing a concise summary of essential patient information. The dashboard also includes a detailed wound view (Fig. 5, right, translated), which

displays the selected wound images alongside the AI-predicted wound area (①). Additionally, the dashboard features trajectory curves for all patient-reported wound scores and AI-derived wound size progression (②). To support continuous model refinement, the system incorporates a feedback mechanism, allowing clinicians to assess the segmentation masks by a simple approval process (③).

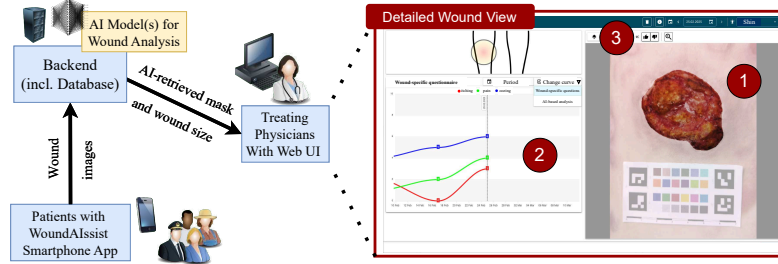


Fig. 5: Simplified visualization of the overall telehealth system.

The system is currently deployed on a secure, GDPR-compliant cloud server and is actively used as part of an ongoing longitudinal clinical study. Approximately 30 patients are expected to participate, each using the *WoundAssist* app on their personal smartphones over six months. This long-term clinical trial aims to empirically assess the effectiveness of the overall telehealth solution, with a focus on patient adherence and wound care outcomes. It also enables the collection of real-world images captured by patients, which provide a valuable basis for both validating the performance of the proposed AI-based wound size estimation and informing future improvements to enhance its robustness.

4 Discussion and Outlook

This study provides a comprehensive benchmark of various DL models for wound segmentation, addressing real-world challenges in clinical deployment and telehealth applications. Interestingly, many of the top-performing models achieved comparable performance on the evaluated datasets, suggesting that binary wound segmentation may be a relatively straightforward task for advanced deep learning methods. This phenomenon may be attributed to the distinctive visual characteristics exhibited by wounds in conventional samples, which enable modern architectures to extract relevant features effectively.

Although no clear trend emerged between CNN- and ViT-based models, the results on both the in-distribution CFU and the DFUC’22 live leaderboard indicate the superior performance of modern GP architectures, such as TransNeXt, over specialized medical models. Similarly, the top five models on the OOD dataset were GP architectures, further emphasizing the superiority of contemporary SS and vision models over specialized medical techniques in handling

OOD data. This enhanced performance is further supported by visual analysis of Grad-CAM heatmaps. Despite the higher computational demands of the top-performing models, inference times remained within practical limits. With a minimum throughput of at least one image per second on the CPU, all models are considered suitable for clinical use and remote wound size monitoring, given the manageable data volumes. Expert evaluations of mask quality and automated size estimates from the top five models yielded promising results, with high mask approval rates and size estimates closely aligning with dermatologist annotations. While no architecture clearly outperformed the others in size retrieval, VW-Conv was most favored by experts for best mask selection.

While our wound size estimation framework, *WoundAmbit*, holds great potential, it also presents limitations and highlights directions for future research:

RO Placement and Robustness Analysis Using Real-World Data.

The evaluation primarily used publicly available and OOD images captured in controlled hospital settings. To improve model robustness and assess size estimation in telemedicine contexts, it is essential to include images from home environments, particularly those taken by patients using smartphones. These images exhibit greater variability in lighting, background, and image quality, which may introduce practical challenges, such as suboptimal RO placement or partial occlusion of ArUco markers. Accurate size estimation requires the RO to be approximately coplanar with the wound surface; however, the current system lacks a mechanism to guide patients in placing the RO correctly. Incorrect positioning can cause perspective distortion and reduced measurement accuracy. Future versions could incorporate real-time user guidance, such as augmented reality feedback, to assist patients during image capture. Moreover, our evaluation included only a single image with partially occluded ArUco markers, limiting our ability to quantify the impact of incomplete or failed marker detection. Future work should involve systematic stress testing to assess how partial occlusions, non-coplanar RO placement, or complete detection failure affect size estimation accuracy. To improve robustness in such scenarios, alternative strategies should be explored to maintain functionality. The patient-captured images from our ongoing longitudinal clinical study can provide a diverse and realistic dataset to evaluate and improve model robustness under suboptimal, real-world conditions.

Model Selection. While efforts were made to select a representative set of models, the evaluation is constrained by the chosen architectures. The three lowest-performing models – U-Net, MISSFormer, and HiFormer – are notably smaller in terms of parameters, which may introduce a slight bias. However, the general trend favoring the GP SS and GP vision models remains consistent, even regarding the larger medical networks.

Over-Reliance on Technology. From a clinical perspective, *WoundAmbit* serves as a decision support system, not a replacement for physicians. While it provides automated wound size tracking, clinicians continue to review both the original wound image and the AI-generated output, ensuring that clinical decisions are not based solely on AI predictions. Nevertheless, the potential risk of over-reliance on technology should be further explored in future studies.

Clinical Validation. Ultimately, clinical validation of *WoundAmbit* is crucial to confirm its effectiveness in practice beyond the scope of technical evaluation. The ongoing longitudinal study will provide critical insight, and follow-up analyses should explore patient adherence over extended periods as well as the system’s overall impact on wound management outcomes and patient well-being.

5 Conclusion

This work advances automated wound segmentation by benchmarking 12 SotA DL methods from GP vision, medical segmentation, and wound care. We address challenges such as domain variability, computational constraints, and model interpretability. Our evaluation on both in-distribution and OOD datasets demonstrates that modern GP models, especially TransNeXt and VWFormer, exhibit strong generalization and remain computationally feasible for clinical use. While our AI-driven wound size estimation framework, *WoundAmbit*, shows strong clinical promise, future work will focus on expanding the dataset with smartphone images from diverse patient demographics and conditions for model finetuning. Additionally, integrating *WoundAmbit* with existing clinical systems and assessing patient-reported outcomes in clinical trials will be crucial for evaluating its real-world impact. Lastly, incorporating the OOD dataset and prospectively collected smartphone images into the training is expected to improve model performance further. Overall, this work represents an important step toward bridging recent advances in computer vision with real-world wound care, offering the potential to improve remote wound size monitoring and enhance patient care.

Acknowledgments. We thank Luisa Deutzmann for her support in creating the 20-image dataset for size retrieval evaluation and Caroline Glatzel for her contributions to manual size annotation and assessment of AI-generated masks.

Ethics The authors declare no conflicts of interest. Ethical approval for this study was waived by the Ethics Committee of the Medical Faculty of the University of Würzburg. All patients provided informed consent for the publication of their photographs.

Supplementary Materials The trained k-fold models can be accessed via Zenodo at 10.5281/zenodo.15123640. The complete source code is archived on Zenodo at 10.5281/zenodo.15682317 and actively maintained on GitHub. Additional materials are provided in the supplementary document available at 10.5281/zenodo.15673941.

References

1. Awais, M., Naseer, M., Khan, S., Anwer, R.M., Cholakkal, H., Shah, M., Yang, M.H., Khan, F.S.: Foundation models defining a new era in vision: A survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **47**(4) (2025)
2. Azad, R., Kazerouni, A., Heidari, M., Aghdam, E.K., Molaei, A., Jia, Y., Jose, A., Roy, R., Merhof, D.: Advances in medical image analysis with vision transformers: A comprehensive review. *Medical Image Analysis* **91** (2024)

3. Borst, V., Riedmann, A., Dege, T., Müller, K., Schmieder, A., Lugrin, B., Kounev, S.: WoundAIssist: A patient-centered mobile app for AI-assisted wound care with physicians in the loop. *arXiv* (2025), <https://arxiv.org/abs/2506.06104>
4. Chairat, S., Chaichulee, S., Dissaneewate, T., Wangkulangkul, P., Kongpanichakul, L.: AI-assisted assessment of wound tissue with automatic color and measurement calibration on images taken with a smartphone. *Healthcare* **11**(2) (2023)
5. Chino, D.Y., Scabora, L.C., Cazzolato, M.T., Jorge, A.E., Traina-Jr, C., Traina, A.J.: Segmenting skin ulcers and measuring the wound area using deep convolutional networks. *Computer Methods and Programs in Biomedicine* **191** (2020)
6. Dhar, M.K., Zhang, T., Patel, Y., Gopalakrishnan, S., Yu, Z.: FUSegNet: A deep convolutional neural network for foot ulcer segmentation. *Biomedical Signal Processing and Control* **92** (2024)
7. Foltynski, P., Ladyzynski, P.: Internet service for wound area measurement using digital planimetry with adaptive calibration and image segmentation with deep convolutional neural networks. *Biocybernetics and Biomedical Engineering* **43**(1) (2023)
8. Goyal, M., Yap, M.H., Reeves, N.D., Rajbhandari, S., Spragg, J.: Fully convolutional networks for diabetic foot ulcer segmentation. In: *IEEE International Conference on Systems, Man, and Cybernetics* (2017)
9. Guo, M.H., Lu, C.Z., Hou, Q., Liu, Z.N., Cheng, M.M., Hu, S.M.: SegNeXt: Rethinking convolutional attention design for semantic segmentation. In: *NeurIPS* (2022)
10. Heidari, M., Kazerouni, A., Soltany, M., Azad, R., Aghdam, E.K., Cohen-Adad, J., Merhof, D.: HiFormer: Hierarchical multi-scale representations using transformers for medical image segmentation. In: *WACV* (2023)
11. Huang, X., Deng, Z., Li, D., Yuan, X., Fu, Y.: MISSFormer: An effective transformer for 2D medical image segmentation. *IEEE Transactions on Medical Imaging* **42**(5) (2022)
12. imito AG: imitoWound (2025), <https://imito.io/en/imitowound>
13. Jørgensen, L.B., Sørensen, J.A., Jemec, G.B., Yderstræde, K.B.: Methods to assess area and volume of wounds – a systematic review. *International Wound Journal* **13**(4) (2016)
14. Kendrick, C., Cassidy, B., Pappachan, J.M., O’Shea, C., Fernandez, C.J., Chacko, E., Jacob, K., Reeves, N.D., Yap, M.H.: Translating clinical delineation of diabetic foot ulcers into machine interpretable segmentation. In: *Diabetic Foot Ulcers Grand Challenge* (2025)
15. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *ICCV* (2023)
16. Langemo, D., Anderson, J., Hanson, D., Hunter, S., Thompson, P.: Measuring wound length, width, and area: Which technique? *Advances in Skin & Wound Care* **21**(1) (2008)
17. Liao, T.Y., Yang, C.H., Lo, Y.W., Lai, K.Y., Shen, P.H., Lin, Y.L.: HardNet-DFUS: Enhancing backbone and decoder of HardNet-MSEG for diabetic foot ulcer image segmentation. In: *Diabetic Foot Ulcers Grand Challenge* (2023)
18. Liu, X., Wang, C., Li, F., Zhao, X., Zhu, E., Peng, Y.: A framework of wound segmentation based on deep convolutional networks. In: *International Congress on Image and Signal Processing, Biomedical Engineering and Informatics* (2017)
19. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1) (2024)

20. Minaee, S., Boykov, Y., Porikli, F., Plaza, A., Kehtarnavaz, N., Terzopoulos, D.: Image segmentation using deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(7) (2022)
21. Oota, S.R., Rowtula, V., Mohammed, S., Liu, M., Gupta, M.: WSNet: Towards an effective method for wound image segmentation. In: *WACV* (2023)
22. Ramachandram, D., Ramirez-GarciaLuna, J.L., Fraser, R.D., Martínez-Jiménez, M.A., Arriaga-Caballero, J.E., Allport, J., et al.: Fully automated wound tissue segmentation using deep learning on mobile devices: Cohort study. *JMIR mHealth and uHealth* **10**(4) (2022)
23. Rayed, M.E., Islam, S.S., Niha, S.I., Jim, J.R., Kabir, M.M., Mridha, M.: Deep learning for medical image segmentation: State-of-the-art advancements and challenges. *Informatics in Medicine Unlocked* (2024)
24. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *MICCAI* (2015)
25. Sanderson, E., Matuszewski, B.J.: FCN-transformer feature fusion for polyp segmentation. In: *Medical Image Understanding and Analysis* (2022)
26. Shaw, J., Bell, P.M.: Wound measurement in diabetic foot ulceration. In: *Global Perspective on Diabetic Foot Ulcerations*, chap. 5. IntechOpen (2011)
27. Shi, D.: TransNeXt: Robust foveal visual perception for vision transformers. In: *CVPR* (2024)
28. Siddique, N., Paheding, S., Elkin, C.P., Devabhaktuni, V.: U-Net and its variants for medical image segmentation: A review of theory and applications. *IEEE Access* **9** (2021)
29. Thisanake, H., Deshan, C., Chamith, K., Seneviratne, S., Vidanaarachchi, R., Herath, D.: Semantic segmentation using vision transformers: A survey. *Engineering Applications of Artificial Intelligence* **126** (2023)
30. Wang, C., Anisuzzaman, D., Williamson, V., Dhar, M.K., Rostami, B., Niezgoda, J., Gopalakrishnan, S., Yu, Z.: Fully automatic wound segmentation with deep convolutional neural networks. *Scientific Reports* **10**(1) (2020)
31. Wang, C., Mahbod, A., Ellinger, I., Galdran, A., Gopalakrishnan, S., Niezgoda, J., Yu, Z.: FUSeg: The foot ulcer segmentation challenge. *Information* **15**(3) (2024)
32. Wang, L., Pedersen, P.C., Agu, E., Strong, D.M., Tulu, B.: Area determination of diabetic foot ulcer images using a cascaded two-stage SVM-based classification. *IEEE Transactions on Biomedical Engineering* **64**(9) (2017)
33. Wang, W., Dai, J., Chen, Z., Huang, Z., Li, Z., Zhu, X., Hu, X., Lu, T., Lu, L., Li, H., et al.: InternImage: Exploring large-scale vision foundation models with deformable convolutions. In: *CVPR* (2023)
34. Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: *ECCV* (2018)
35. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: SegFormer: Simple and efficient design for semantic segmentation with transformers. In: *NeurIPS* (2021)
36. Yan, H., Wu, M., Zhang, C.: Multi-scale representations by varying window attention for semantic segmentation. In: *ICLR* (2024)
37. Zhang, P., Chen, X., Yin, Z., Zhou, X., Jiang, Q., Zhu, W., Xiang, D., Tang, Y., Shi, F.: Interactive skin wound segmentation based on feature augment networks. *IEEE Journal of Biomedical and Health Informatics* **27**(7) (2023)
38. Zhang, R., Tian, D., Xu, D., Qian, W., Yao, Y.: A survey of wound image analysis using deep learning: Classification, detection, and segmentation. *IEEE Access* **10** (2022)