

# SALVE: A 3D Reconstruction Benchmark of Wounds from Consumer-grade Videos

Remi Chierchia<sup>1,2</sup>, Leo Lebrat<sup>1,2</sup>, David Ahmedt-Aristizabal<sup>1,2</sup>, Olivier Salvado<sup>1</sup>,  
Clinton Fookes<sup>1</sup>, Rodrigo Santa Cruz<sup>1,2</sup>  
Queensland University of Technology<sup>1</sup>, CSIRO Data61<sup>2</sup>  
remi.chierchia@hdr.qut.edu.au

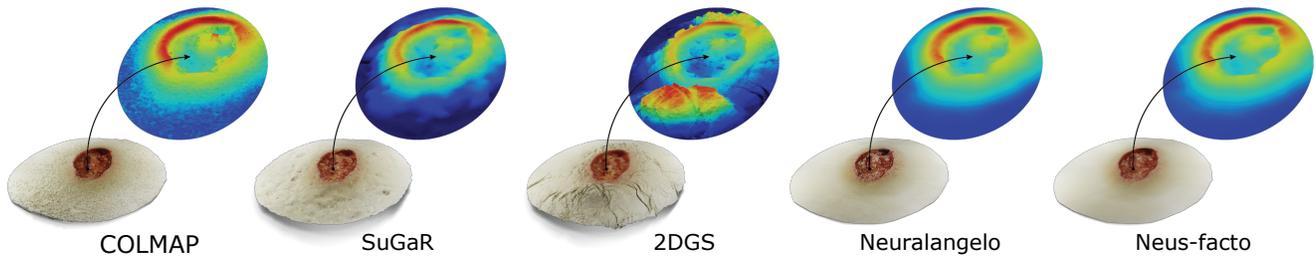


Figure 1. Qualitative evaluation of six 3D surface reconstruction methods using our SALVE dataset. On the left, we present COLMAP, a classic photogrammetric pipeline, which represents typical methods utilised for 3D reconstruction. Next, we display two Gaussian splatting approaches SuGaR [13] and 2DGS [14], followed by two NeRF approaches Neuralangelo [24] and Neus-facto [51]. To highlight surface regularity—important for wound analysis—we present depth color-coded from blue (far regions) to red (closer regions).

## Abstract

Managing chronic wounds is a global challenge that can be alleviated by the adoption of automatic systems for clinical wound assessment from consumer-grade videos. While 2D image analysis approaches are insufficient for handling the 3D features of wounds, existing approaches utilizing 3D reconstruction methods have not been thoroughly evaluated. To address this gap, this paper presents a comprehensive study on 3D wound reconstruction from consumer-grade videos. Specifically, we introduce the SALVE dataset, comprising video recordings of realistic wound phantoms captured with different cameras. Using this dataset, we assess the accuracy and precision of state-of-the-art methods for 3D reconstruction, ranging from traditional photogrammetry pipelines to advanced neural rendering approaches. In our experiments, we observe that photogrammetry approaches do not provide smooth surfaces suitable for precise clinical measurements of wounds. Neural rendering approaches show promise in addressing this issue, advancing the use of this technology in wound care practices. We encourage the readers to visit the project page: SALVE.

## 1. Introduction

Chronic wounds represent a significant health and economic burden worldwide [33, 49]. Effective wound treatments depend on multiple wound clinical measurements, typically performed manually by specialized healthcare professionals [19, 31]. For instance, wound surface area is typically measured by performing planimetry of the wound bed [34]. These procedures are not only invasive and cause patient discomfort but are also prone to errors due to ambiguous definitions of metrics and variations in professionals’ skill levels. Moreover, this method is neither cost-effective nor scalable, as it requires specialized personnel who are often not available in remote areas.

Automatic measurement systems using computer vision and off-the-shelf cameras offer a promising avenue for wound care [5, 12]. However, most existing commercial approaches [1, 2] compute wound measurements solely from 2D images, which are inherently inaccurate for wounds with complex geometry or located on highly curved body parts such as the heel, toe, and lower leg. Also, 2D wound measurements are perspective-dependent. As illustrated in Figure 2, wound measurements such as surface area can vary significantly when computed from images taken from different view angles. Furthermore, these approaches can not estimate wound depth, potentially overlooking a crucial as-

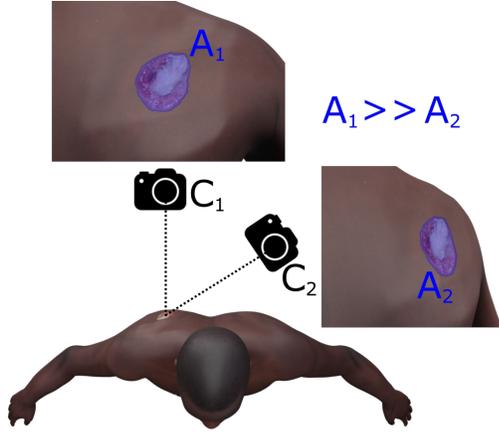


Figure 2. Limitation of 2D Wound Analysis: Images of a wound taken from different view perspectives will present significantly different wound areas.

pect of the wound healing process.

In response to these limitations, 3D reconstruction techniques can be leveraged to generate 3D models from short videos depicting the wound from multiple viewpoints. Subsequently, measurements can be directly computed in 3D. This approach ensures that measurements conducted on the reconstructed geometry are independent of the view direction. Additionally, 3D analysis of wounds allows for the computation of richer wound biomarkers, which could streamline wound documentation [16]. Leveraging these prospects, several works have advocated for 3D wound documentation [4, 6, 9, 15, 16, 21, 22, 26, 40]. However, these studies only considered previous generation 3D reconstruction frameworks, which have recently been surpassed by highly optimized photogrammetric toolboxes and recent neural rendering alternatives.

More importantly, no studies to date have thoroughly evaluated and compared the effectiveness of 3D reconstruction methods specifically applied to wound reconstruction from videos. To address this gap, we introduce a new dataset **SALVE**, designed to capture common challenges encountered in clinical settings, such as complex lighting conditions and varied image quality and resolution of realistic silicone wounds. Using this dataset, we evaluate robust photogrammetry pipelines and modern neural rendering approaches for 3D reconstruction. Our work includes a rigorous evaluation protocol that defines metrics and procedures to assess the geometric accuracy and precision of the evaluated reconstruction algorithms.

We believe that this study is crucial to the development of reproducible research in the area of wound analysis and is a valuable contribution to the field of computer vision through the introduction of a new challenging task. Our benchmark highlights novel avenues for the use of neural 3D reconstruction in a medical setting and exhibits

some limitations of current state-of-the-art (SOTA) methods when applied in this difficult scenario. We highlight the robustness of different 3D reconstruction algorithms with respect to the quality of the acquisition device and demonstrate that some approaches show significant degradation in performance on lower-end devices. We aim for this research to pave the way towards the development of affordable wound telehealth solutions and further reduce geographic inequalities in wound care access.

## 2. Related Work

3D wound analysis has been a topic of interest for at least fifteen years [44], and its advantages over 2D methods and manual procedures have been highlighted in multiple clinical studies in wound care [21, 42]. However, to the best of our knowledge, no study proposing a rigorous and unified evaluation of the accuracy and repeatability of state-of-the-art methods for 3D reconstruction applied to wound reconstruction from consumer-grade cameras exists. Most existing works in this field focus on the repeatability and accuracy of specific wound measurements computed from images, pairs of images, or videos [6, 26, 28, 36, 52]. These studies typically use a single device and follow a particular acquisition protocol, without evaluating the 3D reconstruction itself. For instance, [52] conducted an inter-laboratory comparative study between photogrammetric and 3D scanner-based volume estimation of skin ulcers. Similarly, [26, 28] performed experiments comparing wound area measurements using 2D and 3D methods. Lastly, [36] extended this analysis to nine different wound measurements performed by dermatologists and non-experts, using automatic and semi-automatic tools for wound measurements. The findings of these studies are limited to those specific measurements and cannot be generalized (e.g. different wound shapes, different devices, other wound etiology, and acquisition protocols,...).

Similar to our work, [8, 40, 44] evaluated the accuracy and repeatability of 3D reconstruction methods for wound analysis by comparing 3D models produced by these techniques with 3D models obtained with specialized 3D scanners. However, [40, 44] used a single artificial phantom created by a special effects artist as their input dataset. [8] used short video recordings of cutaneous leishmaniasis skin lesions, which are flat-looking and lack significant 3D features such as depth and complex shapes. These works also focused solely on a single acquisition device and used very simple photogrammetric pipelines. In contrast, our evaluation includes three types of wounds with rich 3D information for proper wound documentation, two acquisition devices, and state-of-the-art photogrammetric pipelines, as well as modern neural rendering approaches for 3D reconstruction. Additionally, unlike existing works, we will make our data and code publicly available for research purposes.



such as COLMAP [38,39] and OpenSfM [32], respectively. However, they rely on accurate feature correspondences between images, which can be challenging in videos recorded in clinical settings due to complex lighting, moving shadows, and self-occlusions. These factors often lead to errors and discontinuities in the reconstructed 3D model.

A novel trend in 3D reconstruction involves using learnable implicit primitives paired [27] with volumetric rendering [17]. As detailed in [27], these approaches express the color of a pixel  $\hat{C}$  as,

$$\hat{C}(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) c(\mathbf{r}(t), \mathbf{d}) dt, \quad (1)$$

where  $\mathbf{r}(t)$  is a time-parametrized ray for the pixel  $\hat{C}$ ,  $c$  is a learnable function depending on the 3D space location and the view direction  $\mathbf{d}$ , responsible for the color information of the neural radiance field, and  $\sigma$  depends only on the spatial coordinate. In this formulation, the density  $\sigma(\mathbf{x})$  can be interpreted as the probability of a ray being stopped at  $\mathbf{x}$ , and a 3D model can be obtained via iso-surface extraction algorithms. However,  $\sigma$  is often implemented as an over-parameterized neural network without surface regularity constraints, which frequently results in 3D reconstructions with floaters or non-smooth surfaces [48]. These surface artifacts could compromise the reliability of clinical wound measurements, making them unsuitable for longitudinal tracking.

To improve the quality of the reconstructed geometry, multiple works [45, 47] propose deriving the density function  $\sigma$  from a learnable Signed Distance Field (SDF), enabling surface regularity modelling within the standard volumetric rendering framework. For instance, in [45],  $\sigma$  is defined to be maximal at the zero-crossings of the SDF  $f_\theta$ ,

$$\sigma(\mathbf{x}) = \max \left( \frac{-H'(f_\theta(\mathbf{x}))}{H(f_\theta(\mathbf{x}))}, 0 \right), \quad (2)$$

where  $H$  is the cumulative distribution function of the logistic distribution, and  $\theta$  represents the model parameters to be learned through volumetric rendering. In this formulation,  $f_\theta$  is subject to the eikonal constraint, and further regularizations, such as curvature [24], can be defined on this implicit representation of objects. The training time for such approaches is in the order of several hours, which limits their applicability in some clinical settings. To reduce the training time associated with the discretization of the rendering Equation (1), a novel family of methods has emerged. They propose an unstructured representation of radiance fields. Gaussian Splatting directly optimizes the position and variance of 3D Gaussians [20]. This representation remains volumetric but can be rasterized more efficiently by applying 2D projection and alpha-compositing.

However, Gaussian splatting does not allow for precise surface extraction. SuGaR [13] and 2DGS [14] proposed

two different formulations amenable to surface extraction from a 3D sparse representation. Although similar in concept, they differ in their problem approach. SuGaR enforces additional constraints on the learned Gaussians to more accurately represent the geometry of the scene. More specifically, the Gaussians should have limited overlap with their neighbors, be as opaque as possible, and be as thin as possible in the normal direction of the surface. Under those assumptions, the neural density of the field can be approximated using,

$$\hat{\sigma}(\mathbf{x}) = \exp \left( -\frac{1}{2s_{g^*}^2} \langle \mathbf{x} - \mu_{g^*}, \mathbf{d}_{g^*} \rangle^2 \right), \quad (3)$$

where  $g^*$  is the index of the closest Gaussian to  $\mathbf{x}$ ,  $s_{g^*}$  is its smallest scaling factor, and  $\mathbf{d}_{g^*}$  is the associated direction. By penalizing the 3D Gaussians to achieve such density, one can significantly enhance the geometric information. Finally, by deriving a set of points and their normals from the Gaussian density, one can employ Poisson reconstruction to swiftly extract a surface mesh with improved quality.

On the other hand, 2DGS proposes to remove the extra dimension  $\mathbf{d}_{g^*}$  and replace the 3D Gaussian with a 2D Gaussian, a multiview-consistent primitive more suited for surface extraction. Both of these explicit formulations reduce runtimes compared to their continuous SDF counterpart. However, they may fall short of providing the regularity required for our application.

**Benchmarked methods:** We compare three proven photogrammetry pipelines: COLMAP [37] (CM), COLMAP equipped with LightGlue feature matching [25] (LGCM), and Meshroom [11] (MR). Additionally, leveraging the computed image poses from LGCM, we benchmark recent AI-based methods for 3D reconstruction. These include Neural Signed Field techniques like NeusFacto [51] and NeuralAngelo [24], as well as discrete explicit approaches such as 3D Gaussian Splatting (3DGS) [20] for rendering and SuGaR [13] or 2D Gaussian Splatting [14] for 3D reconstruction. Finally, we experimented with fast methods using Neural Hash encoding such as InstantNGP [30] for rendering and Neus2 [47], its SDF-based variant. We also explored SfM-free approaches like DUST3R [46], an end-to-end stereo 3D reconstruction method.

### 3.3. Experiments Description

For 3D wound reconstruction methods to be valuable in research and clinical settings, they must be both accurate and precise. Accuracy assesses how closely a reconstruction matches the ground truth, while precision (or reliability and repeatability) measures the consistency between repeated reconstructions. This consistency is essential for verifying the robustness to varying input data and ensuring suitability for tracking wound measurements over time.

		Accuracy																	
Wound	Method	iPhone									Logitech								
		AD	$\mathcal{W}_2$	HD	HD <sub>90</sub>	NC	$\mathcal{W}_2$ -NC	t	CPU	GPU	AD	$\mathcal{W}_2$	HD	HD <sub>90</sub>	NC	$\mathcal{W}_2$ -NC	t	CPU	GPU
PIS3	CM	0.305	1.702	2.757	0.624	0.812	0.797	1.26h	10.1GB	4.73GB	0.696	3.106	14.775	2.131	0.745	0.564	31.2m	2.8GB	2GB
	LGCM	0.259	1.632	2.858	0.516	0.830	0.823	1.29h	11GB	14.1GB	0.751	4.470	12.131	2.653	0.739	0.549	31.7m	2.7GB	4.4GB
	*MR	<b>0.230</b>	<b>1.575</b>	<b>1.272</b>	<b>0.454</b>	<b>0.992</b>	<b>0.991</b>	<b>15.3m</b>	8GB	<b>1.3GB</b>	<b>0.249</b>	<b>1.572</b>	79.029	<b>0.504</b>	<b>0.990</b>	<b>0.990</b>	<b>4.3m</b>	4.7GB	<b>0.7GB</b>
	SuGaR	0.806	4.011	9.657	1.746	0.906	0.822	1.23h	<b>5.7GB</b>	71.1GB	1.075	2.336	6.317	2.342	0.898	0.869	45.6m	6GB	20.4GB
	2DGS	0.757	3.069	6.730	1.695	0.726	0.682	1.66h	23.5GB	31.4GB	1.402	3.826	7.380	3.170	0.789	0.759	27.6m	19.6GB	7.8GB
	Neuralangelo	0.257	1.614	1.359	0.506	0.992	0.992	15.54h	7.1GB	9.7GB	0.982	1.990	14.605	2.123	0.977	0.967	8.17h	4GB	9.7GB
	Neus-facto	0.282	1.669	1.681	0.569	<b>0.993</b>	<b>0.993</b>	55.4m	19GB	24.3GB	<b>0.263</b>	<b>1.612</b>	<b>1.632</b>	<b>0.621</b>	<b>0.993</b>	<b>0.993</b>	46.8m	12.2GB	13.9GB
PIS4	CM	0.215	1.572	2.689	0.499	0.853	0.843	1.27h	10.3GB	4.7GB	0.827	3.668	12.438	2.262	0.806	0.697	31m	2.8GB	2GB
	LGCM	0.216	1.766	2.313	0.503	0.868	0.859	1.3h	10.5GB	14.2GB	1.068	6.638	19.263	3.874	0.766	0.559	32.7m	3.5GB	4.4GB
	*MR	0.192	1.556	2.664	0.475	0.991	0.984	<b>15.8m</b>	7.7GB	<b>1.3GB</b>	0.424	1.779	3.503	1.030	0.982	0.980	<b>4.4m</b>	4.7GB	<b>0.7GB</b>
	SuGaR	0.657	2.003	4.608	1.437	0.950	0.942	1.2h	<b>5.1GB</b>	61.6GB	3.342	7.144	17.882	10.957	0.709	0.462	42.38m	6GB	21.5GB
	2DGS	0.747	2.570	6.364	1.835	0.833	0.808	1.61h	23.5GB	24.2GB	1.057	2.623	7.203	2.409	0.858	0.846	27.3m	16.2GB	5GB
	Neuralangelo	<b>0.178</b>	<b>1.546</b>	2.472	<b>0.427</b>	<b>0.992</b>	<b>0.985</b>	16.37h	7.7GB	9.7GB	0.953	2.076	3.165	2.006	0.972	0.975	8.15h	4.3GB	9.7GB
	Neus-facto	0.217	1.555	<b>2.301</b>	0.539	0.991	<b>0.985</b>	51.8m	19GB	24.2GB	<b>0.256</b>	<b>1.607</b>	<b>2.567</b>	<b>0.639</b>	<b>0.991</b>	<b>0.985</b>	47.6m	12.2GB	14GB
SD	CM	0.355	7.547	9.956	0.730	0.887	0.817	1.2h	10GB	4.7GB	0.963	8.467	26.819	2.962	0.775	0.578	32.7m	3GB	2GB
	LGCM	0.366	7.297	16.488	0.722	0.892	0.819	1.3h	10.5GB	14.1GB	0.595	7.639	23.129	1.648	0.828	0.687	33.7m	3GB	4.4GB
	*MR	0.357	7.765	3.292	0.836	0.979	0.924	<b>13.9m</b>	7.2GB	<b>1.3GB</b>	0.736	7.701	5.767	1.866	0.956	0.933	<b>4.5m</b>	4.7GB	<b>0.7GB</b>
	SuGaR	0.698	8.455	9.625	1.559	0.932	0.823	1.25h	<b>5.6GB</b>	94GB	2.295	11.394	15.475	6.473	0.767	0.597	43.9m	5.9GB	19.5GB
	2DGS	1.195	8.315	8.274	2.484	0.872	0.830	1.65h	23.5GB	24.6GB	1.846	9.076	10.209	3.818	0.849	0.804	26.2m	13GB	5.5GB
	Neuralangelo	<b>0.201</b>	<b>7.149</b>	3.115	<b>0.439</b>	0.980	0.906	14.23h	<b>5.6GB</b>	9.7GB	<b>0.547</b>	7.460	5.529	1.366	0.957	0.935	8.1h	5GB	9.7GB
	Neus-facto	0.230	7.257	<b>3.048</b>	0.510	<b>0.986</b>	<b>0.925</b>	52.5m	19GB	24.2GB	<b>0.428</b>	<b>7.192</b>	<b>5.045</b>	<b>1.034</b>	<b>0.975</b>	<b>0.939</b>	46.2m	12.2GB	14GB

Table 1. Table presenting the accuracy of different 3D reconstruction methods across three wound types (PIS3, PIS4, and SD) and two acquisition devices (iPhone and Logitech). Results highlight the superiority of Meshroom (MR), Neuralangelo, and Neus-facto across six metrics evaluated: AD,  $\mathcal{W}_2$ , HD, and HD<sub>90</sub> (reported in millimetres), along with NC and  $\mathcal{W}_2$ -NC values (ranging from 0 for orthogonal to 1 for the same orientation). Meshroom emerges for best time complexity (t), while SuGaR and 2DGS require the most memory.\*Due to compatibility issues during the setup of Meshroom with our H100 GPUs we used an HPC node with an NVIDIA RTX A6000 instead.

Additionally, these methods can serve as powerful tools for wound visualization. We evaluate the approaches above based on these three criteria, which are detailed below.

**Accuracy Evaluation:** From each recorded video in our dataset, we initially curated a set of fifty frames using the frame selection procedure detailed in Section 3.1. We chose fifty frames as we observed that adding more frames did not enhance wound visibility and sometimes compromised the accuracy of SfM pose estimation due to the inclusion of blurry frames. Using these frames, we generated 3D models of the wounds using the methods outlined in Section 3.2 and compared their results to the ground truth obtained with the specialized 3D scanner detailed in Section 3.1. The outcomes of this experiment are presented in Table 1.

**Precision Evaluation:** We evaluate the consistency of 3D reconstruction methods across different camera devices (inter-device) and recording attempts (inter-recording). From each video in our dataset, we selected 150 sharp frames using the procedure outlined in Section 3.1. These frames were then temporally divided into fifty subsequences, each containing three frames. Within each subsequence, the frames were randomly assigned to separate splits, resulting in three data splits of fifty frames for each video.

For inter-device precision, we compare all pairs of 3D reconstructions generated by each method using data splits generated from different recording devices for the same wound. To assess inter-recording precision, we focus only on data splits derived from videos recorded with the iPhone. Then, we compare all pairs of 3D reconstructions produced

by each method from these data splits of the same wound. The results of this experiment are presented in Table 2.

**Rendering Quality:** To evaluate the suitability of the 3D reconstruction methods as visualization tools (excluding the photogrammetric pipelines), we assessed their rendering capability for view perspectives not present in the input training images. To this end, we extract 75 frames from each video in our dataset using our frame selection procedure, divide them into 50 training images and 25 test images with different poses and visual content, as suggested in [50]. We then train each AI-based 3D reconstruction model on the training set and evaluate their rendering quality on the test images. Table 3 presents the results of this experiment.

**Evaluation Protocol:** The evaluation protocol involves an initial coarse alignment of the 3D reconstructions with the ground-truth acquired data. Specifically, ArUco markers are placed in each scene to retrieve a global transformation and scale the 3D reconstruction for evaluation. We emphasize that these markers are not utilized for pose estimation or any other process by the evaluated methods. A common next step is applying the Iterative Closest Point algorithm (ICP) [35] for fine alignment. However, the convergence of the ICP algorithm can be affected by errors in the reconstructed geometry. Some variations of the algorithm have been proposed to address these issues [35], but these heuristics are often difficult to implement robustly across all evaluations. Consequently, we opted to conduct manual alignment using 3D point cloud and mesh processing software. To ensure a fair and transparent evaluation, we provide the manually aligned meshes alongside the datasets.

Table 2. Table presenting the precision of the best methods from the accuracy experiment. We report precision for inter-recording in the iPhone acquisitions and inter-device across the three wound types. Results highlight Neus-facto as the most robust method among those evaluated.

Wound	Method	Precision											
		Inter recording (iPhone)						Inter device					
		AD	$\mathcal{W}_2$	HD	HD <sub>90</sub>	NC	$\mathcal{W}_2$ -NC	AD	$\mathcal{W}_2$	HD	HD <sub>90</sub>	NC	$\mathcal{W}_2$ -NC
PIS3	MR	<b>0.089</b>	0.671	2.099	<b>0.138</b>	0.994	0.970	0.224	0.617	2.266	0.441	0.991	0.979
	Neuralangelo	0.102	<b>0.405</b>	<b>0.452</b>	0.183	<b>0.997</b>	<b>0.997</b>	0.865	1.578	9.191	1.611	0.981	0.966
	Neus-facto	0.126	0.457	0.476	0.229	<b>0.997</b>	<b>0.997</b>	<b>0.197</b>	<b>0.493</b>	<b>0.854</b>	<b>0.382</b>	<b>0.998</b>	<b>0.998</b>
PIS4	MR	0.216	0.542	1.990	0.499	0.988	0.989	0.441	0.736	2.646	0.946	0.984	0.984
	Neuralangelo	<b>0.093</b>	<b>0.443</b>	1.382	<b>0.160</b>	0.994	0.992	0.645	1.000	3.256	1.482	0.983	0.982
	Neus-facto	0.099	0.449	<b>0.593</b>	0.176	<b>0.995</b>	<b>0.994</b>	<b>0.286</b>	<b>0.561</b>	<b>1.556</b>	<b>0.548</b>	<b>0.992</b>	<b>0.992</b>
SD	MR	0.347	<b>0.826</b>	5.954	0.710	0.978	0.976	0.663	1.676	5.041	1.449	0.961	0.952
	Neuralangelo	<b>0.279</b>	0.888	<b>2.078</b>	<b>0.533</b>	0.984	0.974	<b>0.330</b>	0.916	<b>2.363</b>	<b>0.657</b>	0.979	0.973
	Neus-facto	<b>0.279</b>	0.935	2.122	0.608	<b>0.986</b>	<b>0.987</b>	0.346	<b>0.886</b>	2.516	0.709	0.984	<b>0.984</b>

**Metrics:** To ensure that our evaluation concentrates solely on the wound area, we define a polygon around the wound in both the reconstructed mesh and the ground-truth point cloud. Within this polygon, we compute the following 3D reconstruction metrics: Average Distance (AD), Earth mover’s Distance ( $\mathcal{W}_2$ ), Hausdorff Distance (at percentiles 100<sup>th</sup> (HD) and 90<sup>th</sup> (HD<sub>90</sub>)), and Normal Consistency (NC). We also leverage the optimal transport couplings computed for  $\mathcal{W}_2$  to measure the normal consistency metric under optimal transport’s assignment ( $\mathcal{W}_2$ -NC). The metrics AD,  $\mathcal{W}_2$ , and HD variants are measured in millimeters, while NC and  $\mathcal{W}_2$ -NC assess the alignment between reconstructed surface normals and ground-truth normals, achieving a maximum value of one for perfect alignment. Finally, for assessing rendering quality, we adopt standard metrics from the NeRF literature: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS).

Additionally, we report in Table 1 reconstruction time (t), CPU and GPU memory usage of each method. Our experiments were conducted on an HPC node equipped with 16 Intel Xeon Platinum 8452Y CPUs, an NVIDIA H100 GPU, and 32GB of RAM.

## 4. Results & Discussion

### 4.1. Accuracy Evaluation

As shown in Table 1, Meshroom, Neuralangelo, and Neus-facto consistently outperform other methods in the accuracy evaluation across various metrics, wound types, and recording devices. In particular, Neuralangelo and Neus-facto excel with complex wound geometries such as SD and PIS4. Additionally, Neus-facto exhibits robustness to variations in recording device quality. Even with Logitech videos more prone to blurriness and with lower resolution compared to iPhone videos, Neus-facto consistently produces good results. It particularly excels in NC and  $\mathcal{W}_2$ -NC, highlighting the benefit of the SDF formulation in generating

smooth reconstructed surfaces.

In contrast, Gaussian Splatting methods struggle with sparsely observed regions, inheriting this limitation from the original 3DGS work. Given our dataset’s realistic scenario design, these methods failed to achieve sufficient accuracy, a finding also reflected in the qualitative assessment of novel-view renderings.

Regarding fast reconstruction methods such as InstantNGP, NeuS2, and DUST3R, the generated meshes exhibit poor quality, containing numerous floaters and defects that do not meet our application’s requirements. Therefore, we have omitted their results from the main manuscript, but we provide a qualitative discussion of their results in the supplementary materials. Additionally, the rendering results of InstantNGP are reported in Table 3.

### 4.2. Precision Evaluation

In Table 2, we examine the repeatability of the reconstructions for the most accurate methods identified in the previous experiment (*i.e.*, Meshroom, Neuralangelo, and Neus-facto) between different camera devices (inter-device column) and recording attempts (inter-recording column). Both Neuralangelo and Neus-facto demonstrate advantages over Meshroom (the most accurate photogrammetry method in our evaluation). Remarkably, Neus-facto exhibits the best performance in the inter-device evaluation, showing the most consistent reconstructions among the compared methods.

### 4.3. Qualitative Analysis

We conduct a qualitative analysis of the evaluated models using error color-coded surfaces. In Figure 4, it appears that photogrammetric pipelines often yield noisy surfaces, primarily due to challenges in stereo matching under uncontrolled illumination in the videos. Similarly, Gaussian splatting approaches encounter similar difficulties due to the Gaussians’ initialization based on the SfM point cloud. While these points reflect the challenges associ-

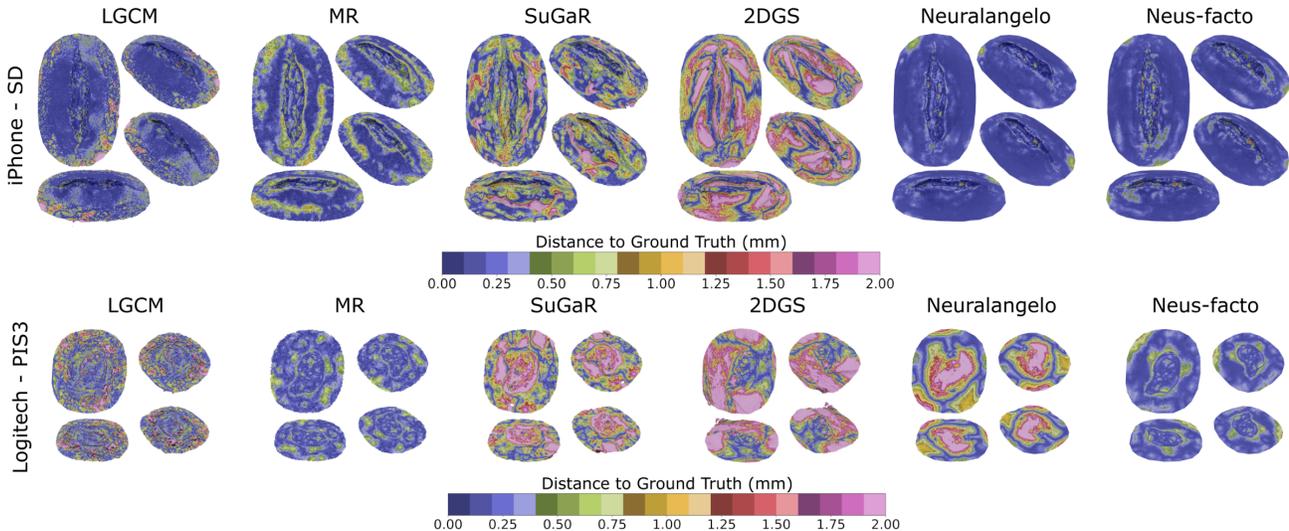


Figure 4. Figure depicting the reconstructed surfaces, where the color represents the distance to the closest point in the ground-truth point cloud, obtained from iPhone and Logitech recordings of SD and PIS3 respectively. Additional results for different wounds and devices are available in the supplementary materials.

Table 3. Table presenting the averaged image metrics for the novel-view renderings experiment. We compare NeRF and Gaussian splatting methods for the iPhone acquisitions of PIS3, PIS4, and SD.

Method	PSNR $\uparrow$			SSIM $\uparrow$			LPIPS $\downarrow$			Time 4k (Mean)
	PIS3	PIS4	SD	PIS3	PIS4	SD	PIS3	PIS4	SD	
SuGaR	29.719	30.935	30.869	0.959	0.960	0.954	0.056	0.050	0.053	1.234h
2DGS	30.874	31.489	31.852	0.966	0.965	0.959	0.055	0.050	0.050	1.695h
3DGS	30.642	32.595	31.904	0.967	0.967	0.960	0.049	0.045	0.046	57.8m
Instant-NGP	30.386	30.979	31.068	0.962	0.962	0.957	0.065	0.059	0.061	6.8m
Neuralangelo	<b>33.897</b>	<b>35.560</b>	35.533	<b>0.971</b>	<b>0.973</b>	<b>0.967</b>	<b>0.047</b>	0.039	<b>0.038</b>	17.884h
Neus-facto	32.224	35.096	<b>35.552</b>	0.964	0.972	0.966	0.066	<b>0.038</b>	<b>0.038</b>	1.225h

ated with stereo-matching, Gaussian splatting—which relies on explicit modelling of the neural surface representation—becomes significantly affected by noise introduced in the photogrammetry pipelines. In contrast, Neuralangelo and Neus-facto leverage implicit volumetric rendering based on the SDF representation of a regular surface, allowing them to regularize noise that may exist in the training data. This results in accurate, repeatable, and smooth reconstructions.

Reproducible performance across different imaging devices is a desirable property for telehealth applications. As detailed in Table 1 and depicted in Figure 5, NeuralAngelo exhibits a massive performance drop when the quality of the input images is reduced, we attribute this behavior to the model’s large size and the reduced photogrammetric consistency between view acquired with a low-resolution device.

#### 4.4. Rendering Quality

In Table 3, Neuralangelo and Neus-facto significantly outperform Instant-NGP (SOTA approach for novel view rendering) in terms of image rendering quality. This su-

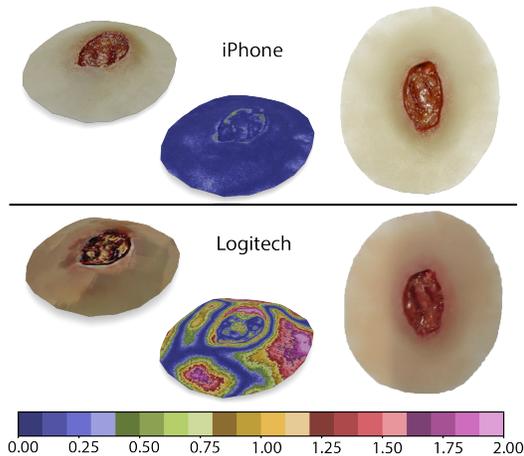


Figure 5. Figure illustrating the influence of image quality on Neuralangelo’s performance. The best reconstruction was obtained with iPhone acquisitions, which feature 4K resolution compared to Logitech’s lower image resolution and quality. From left to right: reconstructed mesh, error color-coded mesh, a sample image of the wound sourced from the training set.

perior performance can be attributed to their ability to accurately represent surfaces, mitigating the floaters commonly found in Instant-NGP and 3DGS, and reducing rendering artifacts present in novel images that are too different from the training distribution (see Figure 6). Despite Instant-NGP’s fast runtime, it was excluded from the quantitative and qualitative evaluations due to its inability to produce surface meshes of good quality. Consequently, Neus-facto emerged as the overall best, balancing the trade-off between

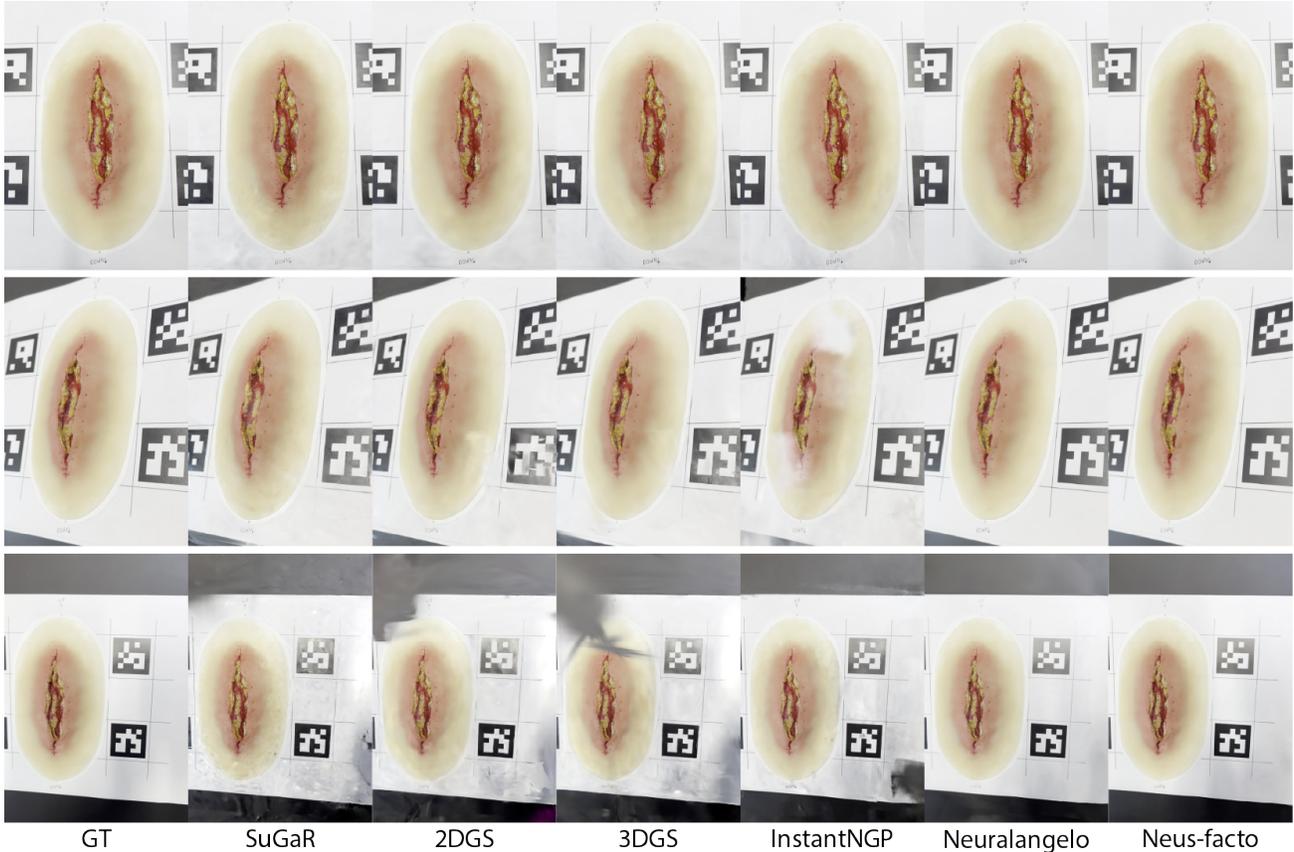


Figure 6. Qualitative evaluation of rendering methods for the SD wound model. The first row depicts an image in the test set that is well-represented in the training set. The second row shows a rendering of the wound from an oblique view, where non-implicit methods may exhibit degraded performance (presence of floaters). The last row displays an image far from the training view distribution, where the rendering can be challenging due to the scarcity of training views.

reconstruction quality, repeatability, and manageable training time. Rendered images and additional qualitative results are made available in the supplementary materials.

## 5. Conclusion

This paper presents a comprehensive study on 3D wound reconstruction using consumer-grade videos. We introduce **SALVE**, a dataset for 3D wound reconstruction and visualization, featuring video recordings of three wound types captured with two different cameras. These videos include photometric challenges typical in clinical settings, such as transient shadows, motion blur, and specular reflections. **SALVE** also provides accurate 3D ground truth acquired with a commercial 3D scanner and a rigorous evaluation protocol for assessing the geometric accuracy and precision of 3D reconstruction methods. Using this dataset, we explore a wide array of 3D reconstruction methods, from state-of-the-art neural rendering to traditional photogrammetry. We believe this study and dataset are crucial for advancing the use of this technology in wound care practices.

Beyond its clinical significance, this study demonstrates some limitations of current AI-based approaches, such as repeatability issues with SfM-initialized methods and the reduced quality of NeuralAngelo’s 3D reconstruction under degraded image quality. We believe that these findings and benchmarks will also serve as valuable resources for the computer vision community.

**Acknowledgments:** This work was supported by the MRFF Rapid Applied Research Translation grant (RARUR000158), CSIRO AI4M Minimising Antimicrobial Resistance Mission, and Australian Government Training Research Program (AGRTP) Scholarship.

**Compliance with Ethical Standards:** This study was performed in line with the principles of the Declaration of Helsinki. The experimental procedures involving human subjects described in this paper were approved by CSIRO Health and Medical Human Research Ethics Committee (CHMHREC) [ethics protocol 2022\_015\_LR].

## References

- [1] imito ag. <https://imito.io/en>. Accessed: 2024-07-12. **1**
- [2] Smartheal. <https://smartheal.org>. Accessed: 2024-07-12. **1**
- [3] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. **3**
- [4] Benjamin Albouy, Emilie Koenig, Sylvie Treuillet, and Yves Lucas. Accurate 3d structure measurements from two uncalibrated views. In *International Conference on Advanced Concepts for Intelligent Vision Systems*, pages 1111–1121. Springer, 2006. **2**
- [5] DM Anisuzzaman, Chuanbo Wang, Behrouz Rostami, Sandeep Gopalakrishnan, Jeffrey Niezgoda, and Zeyun Yu. Image-based artificial intelligence in wound assessment: A systematic review. *Advances in Wound Care*, 11(12):687–709, 2022. **1**
- [6] Fellipe M.C. Barbosa, Bruno M. Carvalho, and Rafael B. Gomes. Accurate chronic wound area measurement using structure from motion. *Proceedings - IEEE Symposium on Computer-Based Medical Systems*, 2020-July:208–213, 7 2020. **2, 3**
- [7] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009. **3**
- [8] Leslie Casas, Sylvie Treuillet, Braulio Valencia, Alejandro Llanos, and Benjamín Castañeda Jr. Low-cost uncalibrated video-based tool for tridimensional reconstruction oriented to assessment of chronic wounds. <https://doi.org/10.1117/12.2070999>, 9287:225–231, 1 2015. **2**
- [9] Damir Filko, Robert Cupec, and Emmanuel Karlo Nyarko. Wound measurement by RGB-D camera. *Machine Vision and Applications*, 29(4):633–654, 5 2018. **2**
- [10] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015. **3**
- [11] Carsten Griwodz, Simone Gasparini, Lilian Calvet, Pierre Gurdjos, Fabien Castan, Benoit Maujean, Gregoire De Lillo, and Yann Lanthony. Alicevision Meshroom: An open-source 3D reconstruction pipeline. In *Proceedings of the 12th ACM Multimedia Systems Conference - MMSys '21*. ACM Press, 2021. **4**
- [12] Giuseppe Guarro, Federico Cozzani, Matteo Rossini, Elena Bonati, and Paolo Del Rio. Wounds morphologic assessment: application and reproducibility of a virtual measuring system, pilot study. *Acta Bio Medica: Atenei Parmensis*, 92(5), 2021. **1**
- [13] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5354–5363, 2024. **1, 4**
- [14] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers, SIGGRAPH '24*, New York, NY, USA, 2024. Association for Computing Machinery. **1, 4**
- [15] Jan Juszczyk, Agata Wijata, Joanna Czajkowska, Marta Biesok, Bartłomiej Pyciński, and Ewa Pietka. Evaluation of Methods for Volume Estimation of Chronic Wounds. *Advances in Intelligent Systems and Computing*, 1011:258–267, 2019. **2**
- [16] Jan Maria Juszczyk, Agata Wijata, Joanna Czajkowska, Michal Krecichwost, Marcin Rudzki, Marta Biesok, Bartłomiej Pyciński, Jakub Majewski, Jacek Kostecki, and Ewa Pietka. Wound 3d geometrical feature estimation using poisson reconstruction. *IEEE Access*, 9:7894–7907, 2020. **2**
- [17] James T Kajiya and Brian P Von Herzen. Ray tracing volume densities. *ACM SIGGRAPH computer graphics*, 18(3):165–174, 1984. **4**
- [18] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)*, 32(3):1–13, 2013. **3**
- [19] David H Keast, C Keith Bowering, A Wayne Evans, Gerald L Mackean, Catherine Burrows, and Lincoln D'souza. Contents: Measure: A proposed assessment framework for developing best practice recommendations for wound assessment. *Wound Repair and Regeneration*, 12:s1–s17, 2004. **1**
- [20] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4), July 2023. **4**
- [21] Syamantak Kumar, Dhruv Jaglan, Nagarajan Ganapathy, and Thomas Martin Deserno. A comparison of open source libraries ready for 3D reconstruction of wounds. 10954:69–76, 3 2019. **2**
- [22] Léo Lebrat, Rodrigo Santa Cruz, Remi Chierchia, Yulia Arzhaeva, Mohammad Ali Armin, Joshua Goldsmith, Jeremy Oorloff, Prithvi Reddy, Chuong Nguyen, Lars Petersson, et al. Syn3DWound: a synthetic dataset for 3d wound bed analysis. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2024. **2, 3**
- [23] Bruno Lévy and Alain Filbois. Geogram: a library for geometric algorithms. 2015. **3**
- [24] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2023. **1, 4**
- [25] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local Feature Matching at Light Speed. In *ICCV*, 2023. **4**
- [26] Chunhui Liu, Xingyu Fan, Zhizhi Guo, Zhongjun Mo, Eric I.Chao Chang, and Yan Xu. Wound area measurement with 3D transformation and smartphone images. *BMC Bioinformatics*, 20(1):1–21, 12 2019. **2**
- [27] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf:

- Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020. 4
- [28] Houman Mirzaalian Dastjerdi, Dominique Töpfer, Stefan J. Rupitsch, and Andreas Maier. Measuring surface area of skin lesions with 2D and 3D algorithms. *International Journal of Biomedical Imaging*, 2019, 2019. 2
- [29] Pierre Moulon, Pascal Monasse, Romuald Perrot, and Renaud Marlet. OpenMVG: Open multiple view geometry. In *International Workshop on Reproducible Research in Pattern Recognition*, pages 60–74. Springer, 2016. 3
- [30] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 4, 1
- [31] ELIZABETH Nichols. Wound assessment part 1: how to measure a wound. *Wound Essentials*, 10(2):51–5, 2015. 1
- [32] OpenSfM. Open source structure-from-motion pipeline (opensfm). 4
- [33] Rosana E Pacella, Ruth Tulleners, Qinglu Cheng, Ellen Burkett, Helen Edwards, Stephen Yelland, David Brain, John Bingley, Peter Lazzarini, Jason Warnock, et al. Solutions to the chronic wounds problem in australia: a call to action. *Wound Practice & Research: Journal of the Australian Wound Management Association*, 26(2):84–98, 2018. 1
- [34] David Reifs, Lorena Casanova-Lozano, Ramon Reig-Bolaño, and Sergi Grau-Carrion. Clinical validation of computer vision and artificial intelligence algorithms for wound measurement and tissue classification in wound care. *Informatics in Medicine Unlocked*, 37:101185, 2023. 1
- [35] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. In *Proceedings third international conference on 3-D digital imaging and modeling*, pages 145–152. IEEE, 2001. 5
- [36] David Sánchez-Jiménez, Fernando F. Buchón-Moragues, Begoña Escutia-Muñoz, and Rafael Botella-Estrada. SfM-3DULC: Reliability of a new 3D wound measurement procedure and its accuracy in projected area. *International Wound Journal*, 19(1):44–51, 1 2022. 2
- [37] Johannes Lutz Schönberger and Jan-Michael Frahm. COLMAP: A general-purpose Structure-from-Motion (SfM) and Multi-View Stereo (MVS) pipeline. 4
- [38] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4
- [39] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 4
- [40] Tim Shirley, Dmitri Presnov, and Andreas Kolb. A lightweight approach to 3D measurement of chronic wounds. *Journal of WSCG*, 27(1):67–74, 2019. 2, 3
- [41] Ashish Sinha, Jeremy Kawahara, Arezou Pakzad, Kumar Abhishek, Matthieu Ruthven, Enjie Ghorbel, Anis Kacem, Djamila Aouada, and Ghassan Hamarneh. Dermsynth3d: Synthesis of in-the-wild annotated dermatology images. *Medical Image Analysis*, 95:103145, 2024. 3
- [42] Ekaterina Sirazitdinova and Thomas M. Deserno. System design for 3D wound imaging using low-cost mobile devices. <https://doi.org/10.1117/12.2254389>, 10138:258–264, 3 2017. 2
- [43] Noah Snavely, Steven M Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. In *ACM siggraph 2006 papers*, pages 835–846. 2006. 3
- [44] Sylvie Treuillet, Benjamin Albouy, and Yves Lucas. Three-dimensional assessment of skin wounds using a standard digital camera. *IEEE Transactions on Medical Imaging*, 28(5):752–762, 5 2009. 2
- [45] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 4
- [46] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 4, 1
- [47] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3295–3306, 2023. 4, 1
- [48] Frederik Warburg, Ethan Weber, Matthew Tancik, Aleksander Holynski, and Angjoo Kanazawa. Nerfbusters: Removing ghostly artifacts from casually captured nerfs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18120–18130, 2023. 4
- [49] Jo Wilkie, Keryln Carville, Shih Ching Fu, Rhonda Kerr, Kathleen Finlayson, Tanya Tuffrey, Jason C Lenzo, and Gary Geelhoed. Determining the actual cost of wound care in australia. *Wound Practice & Research: Journal of the Australian Wound Management Association*, 31(1):7–18, 2023. 1
- [50] Wenhui Xiao, Rodrigo Santa Cruz, David Ahmedt-Aristizabal, Olivier Salvado, Clinton Fookes, and Leo Lebrat. Nerf director: Revisiting view selection in neural volume rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024. 5
- [51] Zehao Yu, Anpei Chen, Bozidar Antic, Songyou Peng, Apratim Bhattacharyya, Michael Niemeyer, Siyu Tang, Torsten Sattler, and Andreas Geiger. Sdfstudio: A unified framework for surface reconstruction, 2022. 1, 4
- [52] Omar Zenteno, Eduardo González, Sylvie Treuillet, Benjamin Castañeda, Braulio Valencia, Alejandro Llanos, and Yves Lucas. Volume estimation of skin ulcers: Can cameras be as accurate as laser scanners? *Lecture Notes in Computational Vision and Biomechanics*, 27:735–744, 2018. 2

# SALVE: A 3D Reconstruction Benchmark of Wounds from Consumer-grade Videos

## Supplementary Material

In this supplementary material, we present further details on the acquisition setup and additional qualitative results, including error color-coded meshes and novel-view renderings. Additionally, we present the non-suitability of the surfaces from DUST3R [46], Instant-NGP [30], and NeuS2 [47] for geometric assessment.

In Section A we present more details relative to the SALVE’s dataset challenges. Next, in Section B we describe more precisely the evaluation protocol, namely a qualitative evaluation of the ground-truth point clouds and specifications about sampling of the meshes adopted for the experiments of Section 3.3 of the main paper. Following, with respect to Figure 4 and Figure 6 of the main paper, we present the additional error color-coded meshes (Figure 4) and renderings (Figure B.4 for PIS3 and Figure B.5 for PIS4) for the remaining devices and wound types we omitted in the main paper.

In addition, we present the surfaces extracted from DUST3R, Instant-NGP, and NeuS2 in Figures B.6, B.7, and B.8 respectively, and justify why those methods are not included in 3D reconstruction benchmark.

Finally, in Figure B.9 we present some cases of failed reconstruction for the photogrammetric approaches when we input a greater amount of images, specifically 100 and 150. Compared to a set of 50 images, 100 and 150 present more blurring artefacts as mentioned in Section 3.1 of the main paper. As a result, we chose sets of 50 images for our SALVE dataset as they provided more consistent results and better wound representations when compared to fewer images.

### A. Dataset additional challenges

As depicted in Figure 3 of the main paper, we use a Logitech 4K webcam and an iPhone 14 Pro Max for video recordings, and a Revopoint POP 3D scanner to acquire the ground-truth point clouds. Both recording devices utilise default acquisition settings (e.g. exposure time, frame rate, aperture, white balance, etc.) to simulate acquisition scenarios as close as possible to a telehealth application, Figure A.1 shows an example of dynamic lighting conditions comprised in our dataset. Furthermore, when extracting frames from the recordings, we select the sharpest frames as outlined in Section 3.1 of the main paper, however, we might still observe blurred images in the datasets. An example regarding the Logitech recordings is partially observable from the image in the second row of Figure A.1. Addition-

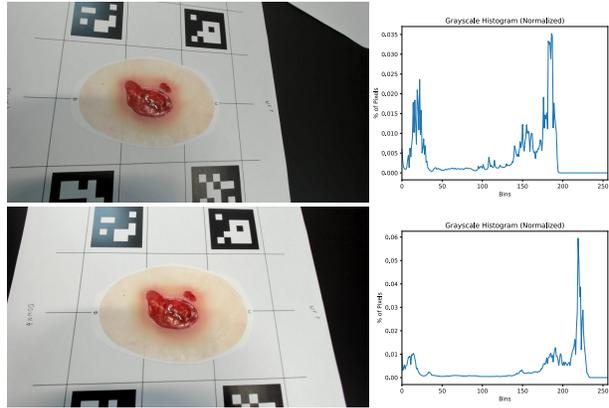


Figure A.1. Two images of the PIS3 wound type captured from slightly different angles using the Logitech camera show different luminosity levels, as indicated by their normalized grayscale histograms (on the right).

ally, SALVE presents moving shadows cast by the operator during the acquisition, which is a common issue in environments with multiple illumination sources.

### B. Evaluation details

In Figure B.2, we illustrate qualitatively the three ground-truth point clouds acquired with the Revopoint POP 3D scanner.

For the metrics AD, HD, HD<sub>90</sub>, and NC we uniformly sampled 2 million points from the reconstructed meshes, while for  $\mathcal{W}_2$  and  $\mathcal{W}_2$ -NC we sampled around 30 thousand points from the reconstructed meshes due to the computational complexity of optimal transport metrics.



Figure B.2. Ground-truth point clouds acquired with the Revopoint POP 3D scanner and their respective total number of points. From left to right: PIS3, PIS4, and SD wound types.

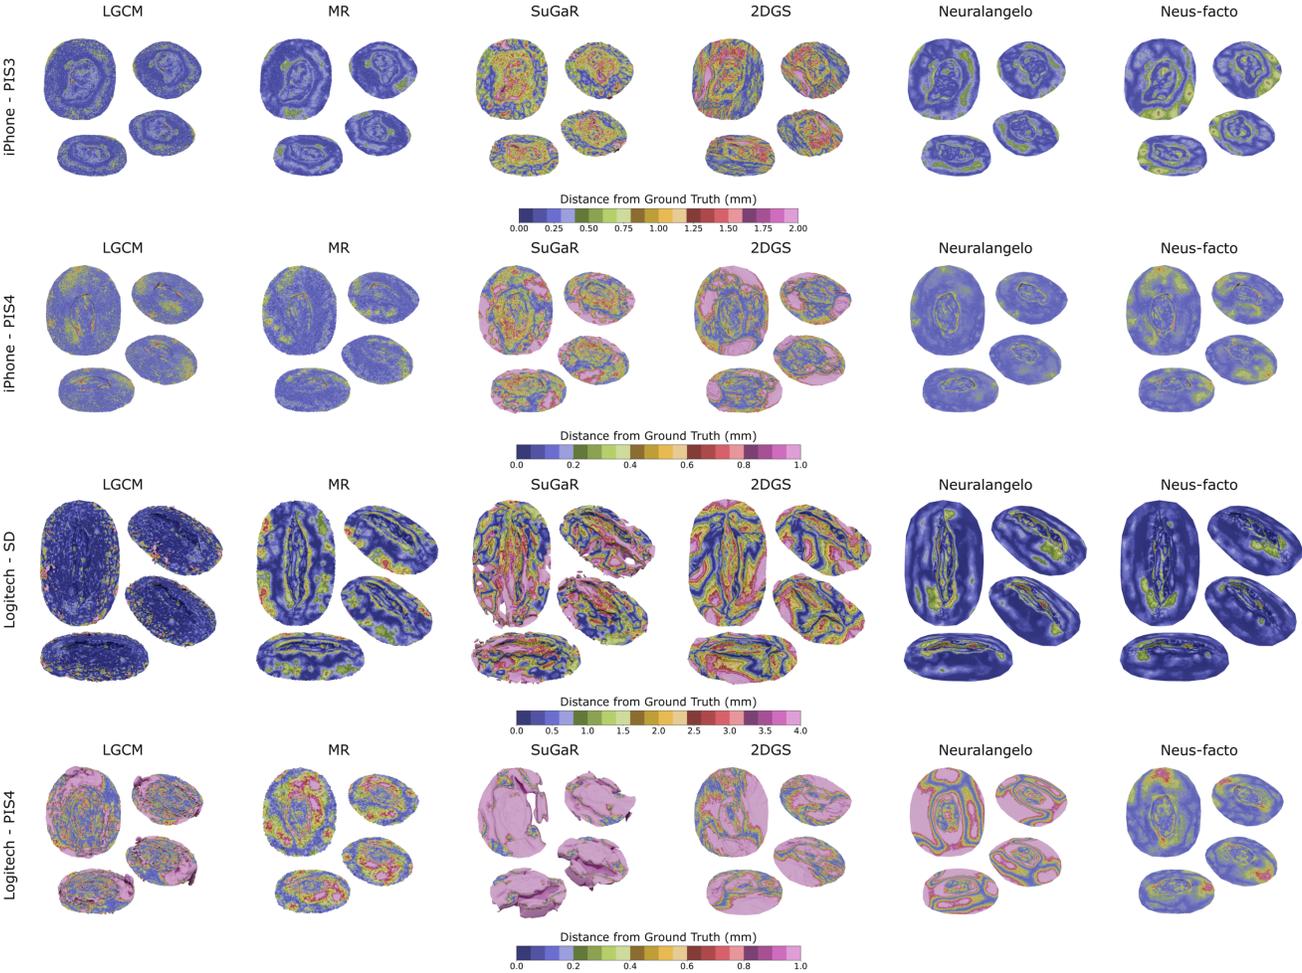


Figure B.3. Reconstructed surface results for different methods, wound types and recording devices. The color represents the distance to the closest point in the ground-truth point cloud.

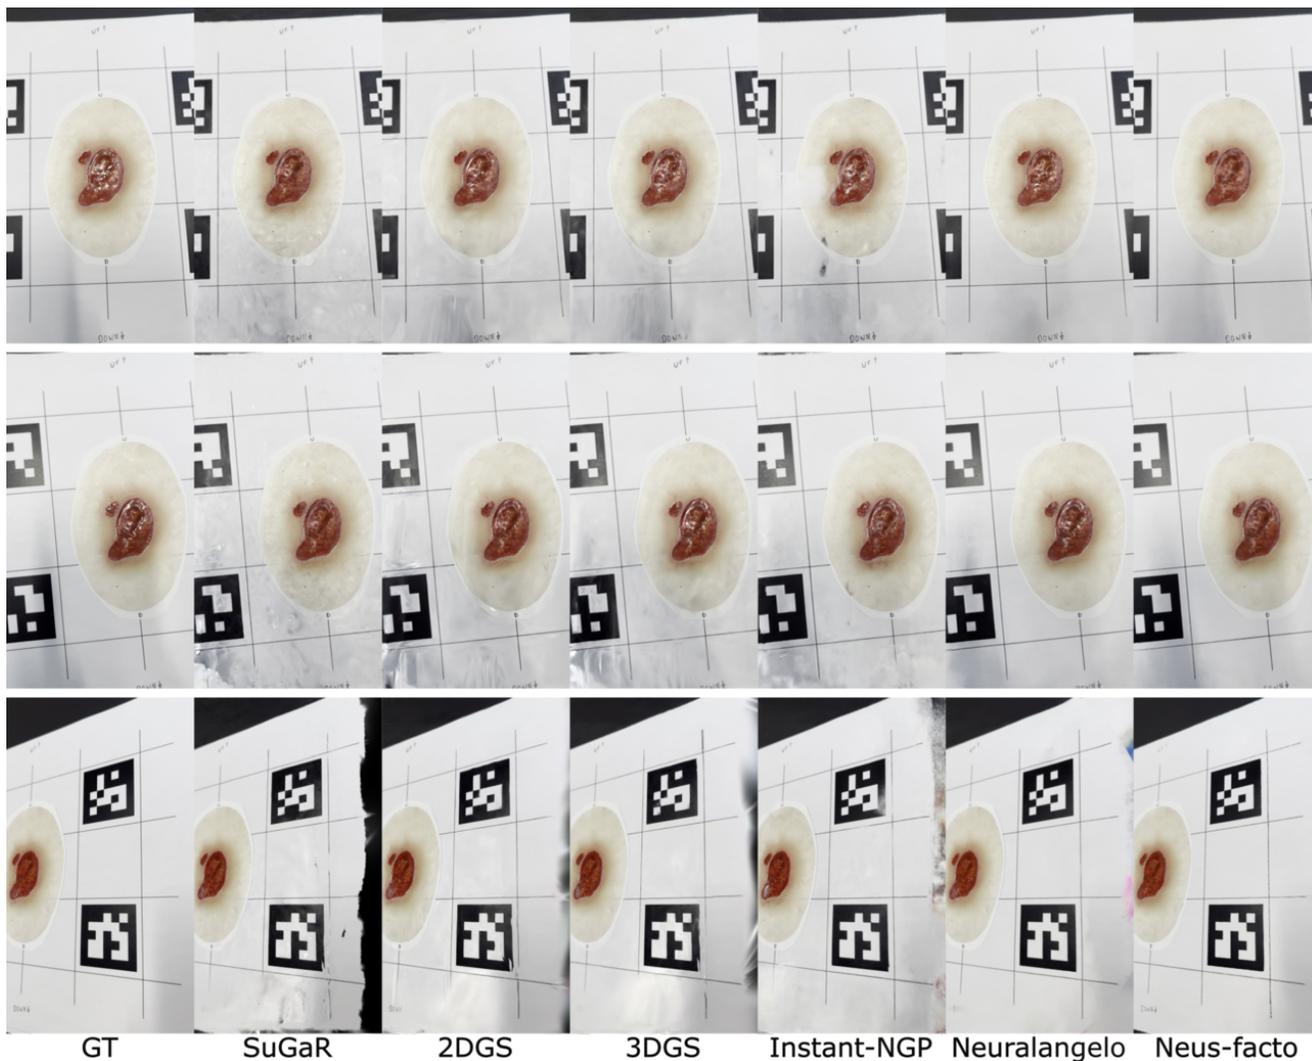


Figure B.4. Qualitative evaluation of rendering methods for the PIS3 wound model. The first row presents the renderings for a view perspective in the test set that is well-represented in the training set. The second row shows a rendering of the wound from an oblique view, where we can observe the presence of floaters in non-SDF-based methods. The last row displays a view perspective under-observed in the training set. Notably, the PIS3 wound type presents hard reflections that appear especially in the first two rows.

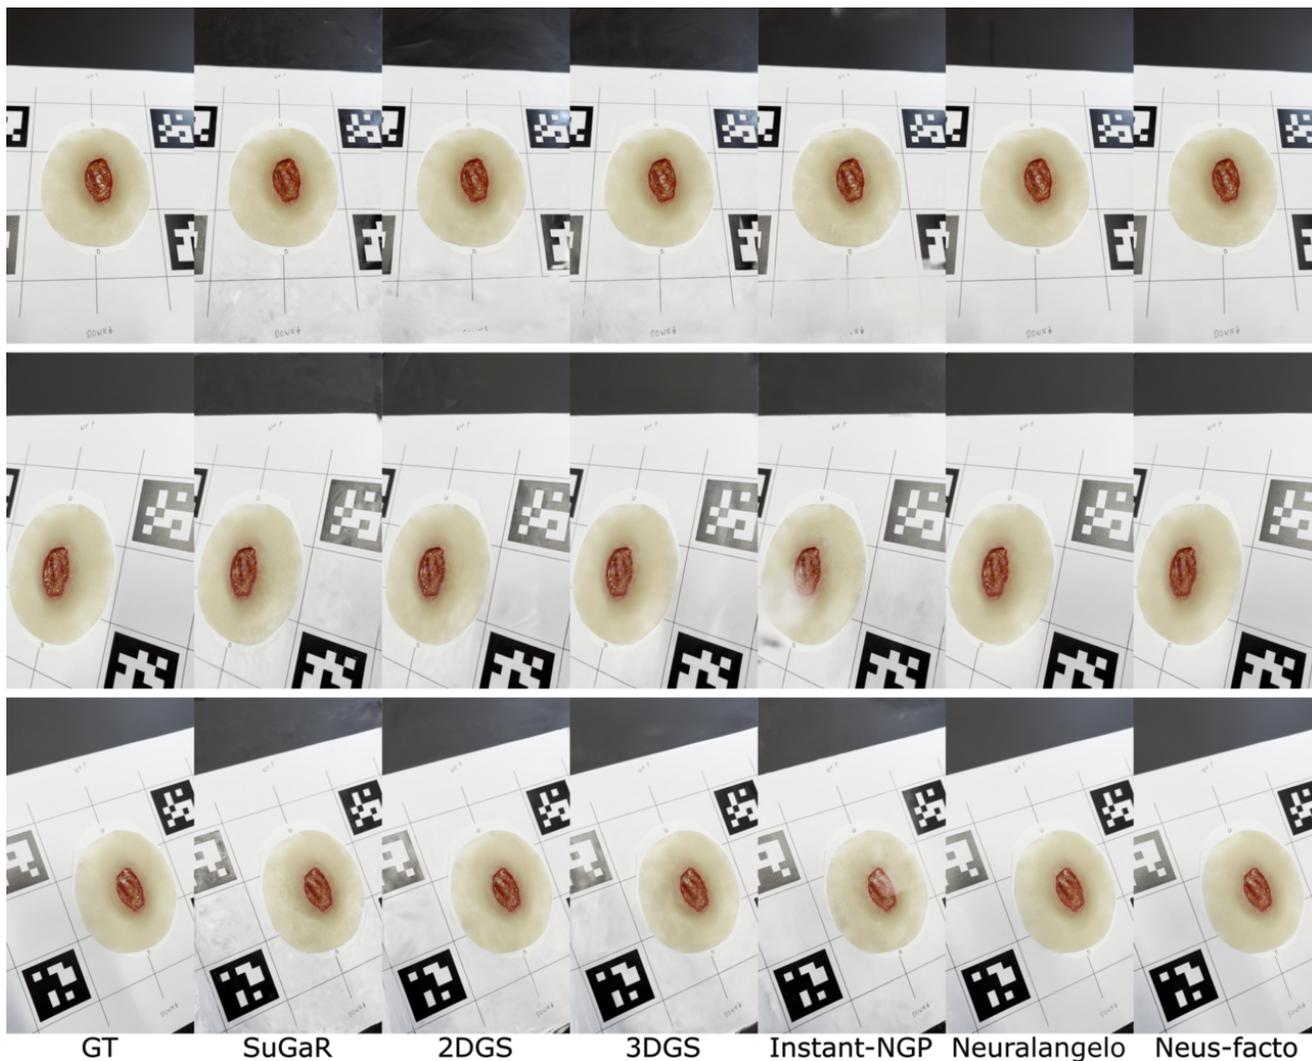


Figure B.5. Qualitative evaluation of rendering methods for the PIS4 wound model. The first row presents the renderings for a view perspective in the test set that is well-represented in training set, however, floaters are already visible for non-SDF-based methods. The second row shows a rendering of the wound from an oblique view, where more floaters are present. The last row displays an image from another view angle where artifacts similar to those in the second row can be observed.

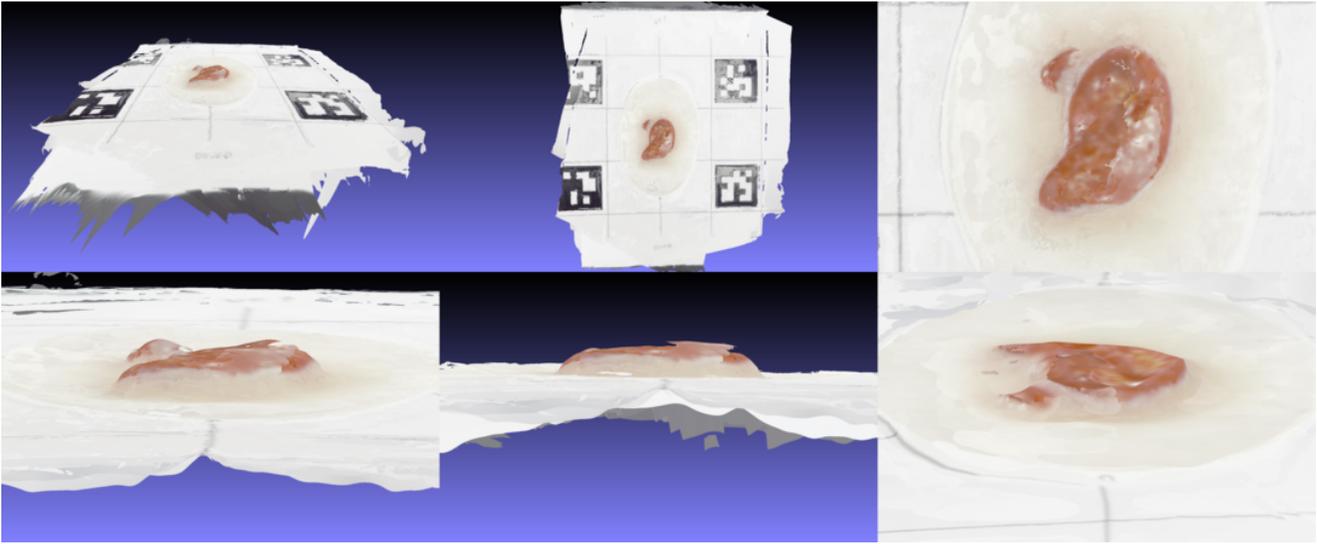


Figure B.6. Six views of the mesh generated by DUST3R for the PIS3 wound type. From the middle figure of the second row, we can observe how DUST3R does not retrieve accurate geometry. The perimeter of the wound is much higher than the surrounding regions, not reflecting the real geometry of the scene. DUST3R, unlike traditional multi-view stereo approaches, do not follow epipolar constraints to generate the 3D structure but relies on transformers tailored to solve reconstructions from a few views in the wild.

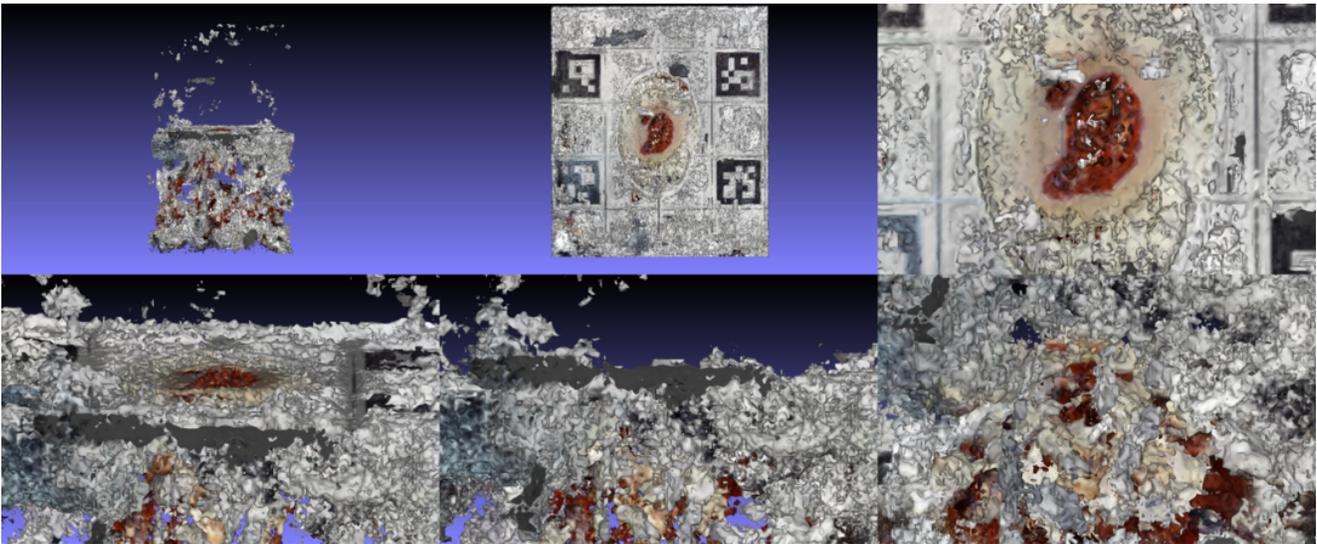


Figure B.7. Six views of the mesh generated by Instant-NGP for the PIS3 wound type. As Instant-NGP is not a method developed for 3D surface reconstruction, its NeRF density is not regularised. As a result, when extracting a mesh using the Marching Cubes algorithm, it presents a structure similar to the one reported in the figure. Given the absence of surface regularizer terms, Instant-NGP solves photogrammetric consistency by placing density values under the wound's real surface, creating a scattered reconstruction.

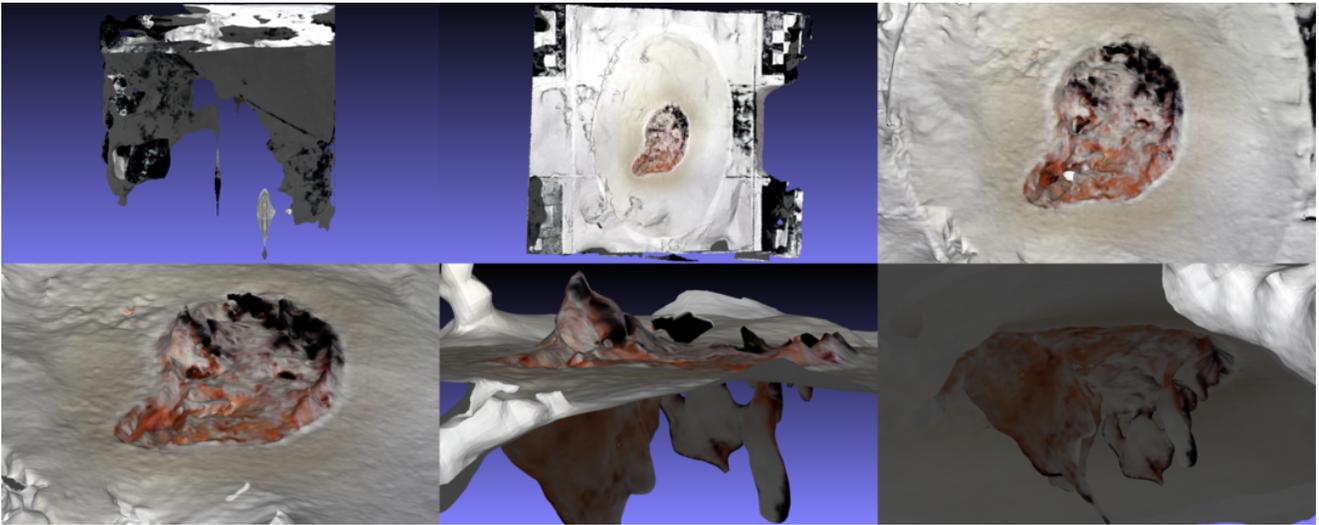


Figure B.8. Six views of the mesh generated by NeuS2 for the PIS3 wound type. NeuS2 is a fast method for 3D surface reconstruction, however, we excluded it from our benchmark because it did not perform consistently well in our **SALVE** dataset. Although more fine-tuning might be required for **SALVE**, we considered Neuralangelo and Neus-facto as better representing SDF methods in terms of robustness against photogrammetric complexities. For example, PIS3 presents hard reflections, as analysed in the renderings evaluation above, and we can observe in both figures on the right how NeuS2 attempts to satisfy photometric consistency and surface regularization by “pushing” the surface under the wound level. This behaviour can be attributed to its architecture similar to Instant-NGP.

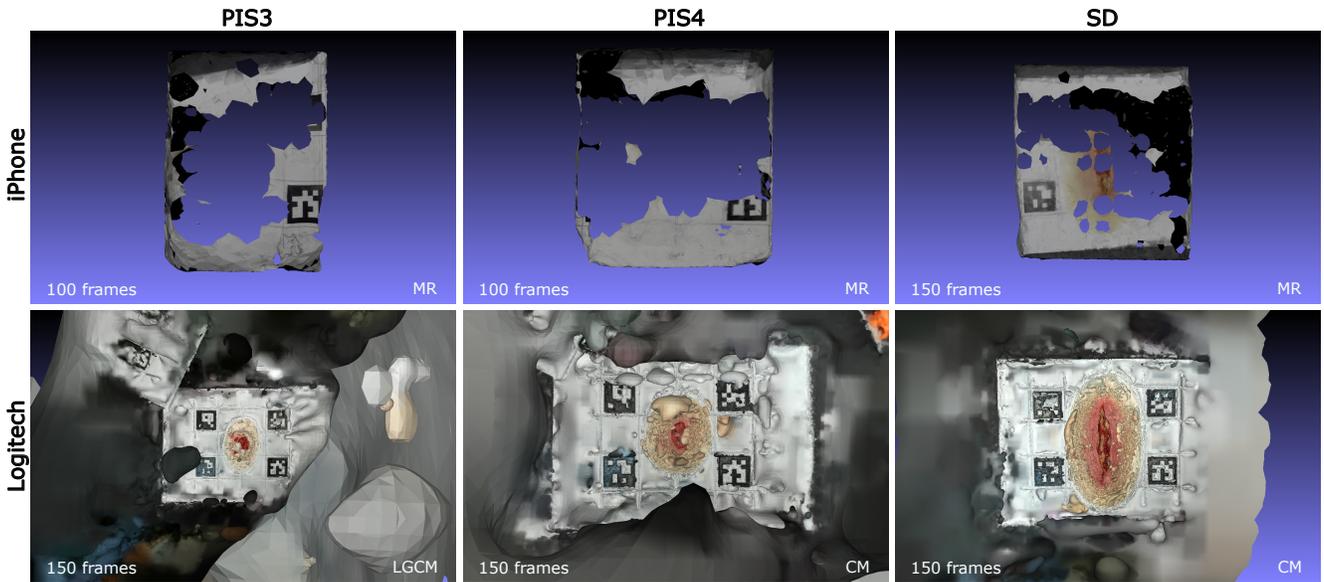


Figure B.9. Examples of failed reconstruction while exploring how larger samples of images impact photogrammetric methods. In the first row, we display examples of 100 and 150 frames sampled from the iPhone sequence of each wound type. Meshroom (MR) was not able to reconstruct any detail in the wound and surrounding regions. In the second row, we display both COLMAP (CM) and COLMAP equipped with LightGlue feature matching (LGCM) at 150 frames. While both methods manage to reconstruct the wound area, the quality of the estimation drastically decreases compared to samples of 50 images.