Privacy-Aware 3D Reconstruction and Obfuscation in the Metaverse

Viviana Pentangelo* Salam Tabet† Stefano Lambiase* Ayman Kayssi† Imad H. Elhajj† Fabio Palomba*
*Department of Computer Science, University of Salerno, Salerno, Italy
†Department of Electrical and Computer Engineering, American University of Beirut, Beirut, Lebanon
vpentangelo@unisa.it, sat22@aub.edu.lb, slambiase@unisa.it, ayman@aub.edu.lb, ie05@aub.edu.lb, fpalomba@unisa.it

***Abstract*—The automated reconstruction of 3D environments is crucial for immersive metaverse experiences, as it enables realistic and dynamic virtual spaces that enhance user interaction and presence; nevertheless, it could raise significant privacy concerns. Image-to-3D techniques, such as photogrammetry and Neural Radiance Fields (NeRF), can inadvertently capture and render sensitive information, posing ethical and legal risks. To address this issue, this work evaluates the combination of obfuscation techniques with 3D reconstruction methods to mitigate privacy threats while maintaining visual quality. Various obfuscation approaches are applied to input images before reconstruction, and their impact on model fidelity and privacy preservation is assessed through quantitative metrics and qualitative discussion. The findings highlight trade-offs between privacy protection and 3D reconstruction quality, where diffusion-based inpainting combined with weaker reconstruction methods can achieve strong privacy preservation for volumetric objects (up to 177% improvement) with lightweight meshes (up to 6.5× less) but slightly lower visual fidelity (less than 5% difference), whereas machine learning-based reconstruction methods can reconstruct challenging surfaces with high realism (up to 80% realism) at the cost of larger meshes (10k–56k polygons) and reduced privacy (56% less on average). This study informs researchers and practitioners on balancing realism and privacy, contributing to the development of privacy-aware metaverse environments.**

***Index Terms*—Metaverse, 3D Reconstruction, Obfuscation Methods, Privacy-Threatening.**

I. Introduction

The development of high-quality 3D environments is a fundamental requirement for constructing immersive metaverse experiences. The realism and detail of virtual spaces directly influence user engagement, making the creation of complex 3D models a critical aspect of metaverse development [9], [19]. However, the manual creation of 3D models is labor-intensive, requiring substantial time and expertise in 3D modeling tools such as Blender and Autodesk. To mitigate these challenges, researchers explored automation techniques that leverage computer vision (CV), particularly image-to-3D methods like photogrammetry and Neural Radiance Fields (NeRF). These techniques allow for the reconstruction of 3D models from 2D images of real objects, significantly reducing the effort and time required for scene creation [1], [23].

Despite the advantages of automation, a major concern in using image-to-3D approaches is the potential capture and rendering of sensitive or privacy-threatening elements [29]. Unlike manually curated models, automated reconstruction techniques can inadvertently include private details, such as faces, license plates, or personal artifacts, which, when rendered in a virtual environment, may lead to ethical and legal issues. Furthermore, privacy-sensitive information is not limited to identifiable faces but extends to personal attributes such as tattoos, religious symbols, and medical items, which are often unintentionally captured during 3D scanning [36]. The immersive nature of the metaverse further exacerbates these concerns, as high-resolution 3D reconstructions can expose private objects and spaces, posing significant privacy risks [22].

Researchers have developed obfuscation techniques to mitigate privacy risks by selectively altering sensitive elements in images and videos before their transformation into 3D models [21]. Traditional methods, such as black masking, blurring, and pixelation [16], [31], have advanced into inpainting techniques that remove objects while preserving scene realism [12]. Recent approaches employ convolutional neural networks (CNNs) [11] and generative adversarial networks (GANs) [14] to improve obfuscation quality. As privacy concerns in the metaverse gain increasing attention, several works started to investigate how to employ them in the metaverse (e.g., [20], [30]). However, to the best of our knowledge, previous research has not empirically tested the integration of obfuscation and 3D reconstruction techniques to generate privacy-preserving 3D models for the metaverse, aiming to balance privacy protection with the visual quality and usability of the models [17], [32]. Advancing our understanding of this integration could yield critical insights, such as identifying optimal trade-offs between privacy and fidelity, assessing the impact on user experience, and informing the development of more robust privacy-preserving frameworks for virtual environments. Unlike applying obfuscation and 3D reconstruction separately, their combined use may mitigate the limitations of each: obfuscation alone can degrade model quality, while reconstruction without privacy controls may expose sensitive details. A well-designed integration could therefore enhance both security and realism, ensuring models remain functional and immersive without compromising personal data.

Building on these considerations, this work takes a step toward the development of a pipeline that ensures both the quality and privacy compliance of 3D models generated from images. The study conducts a comparative analysis of different configurations for (1) removing privacy-threatening elements from input images using obfuscation and (2) generating 3D

models using different 3D reconstruction approaches. The effectiveness of these configurations is evaluated considering both the visual quality of the resulting models and their adherence to privacy-preserving constraints. The findings highlight trade-offs between automation, quality, and privacy, providing insights and guidelines for future steps.

II. Background and Related Work

In this section, we overview the most closely related literature addressing both privacy protection and 3D reconstruction methodologies.

*a) **3D Reconstruction Techniques**:* Aharchi et al. [1] categorized CV methods in active and passive techniques. In our study, we are interested in the latter, primarily photogrammetry, 3D Gaussian-splatting (3DGS), and Neural Radiance Fields (NeRF)-based models. Photogrammetry reconstructs 3D models by analyzing multiple overlapping images taken from different perspectives. This technique extracts depth, scale, and geometric properties through computational methods, producing accurate digital representations of objects and environments [4]. Similarly, 3DGS is a technique that creates 3D models from images or point clouds, representing the scene as a set of Gaussian ellipsoids. These are projected and rasterized to generate images from different angles, allowing for fast and fluid visualization [35]. NeRF-based models have gained significant attention for their ability to generate photorealistic 3D reconstructions. Introduced by Mildenhall et al. [18], NeRF leverages deep learning to synthesize 3D scenes from multiple 2D images. Unlike traditional approaches that generate explicit meshes or point clouds, NeRF learns how light behaves within a scene, producing highly realistic views from arbitrary angles. This is achieved by training a neural network to represent the scene as a continuous function mapping spatial coordinates to color and light density. Later research refined NeRF to enhance practical applicability [26].

*b) **Obfuscation Techniques**:* Ensuring privacy while scanning 3D scenes for the metaverse is crucial, as high-resolution 3D reconstruction can expose sensitive personal information [2], [22]. To address these risks, various obfuscation techniques have been developed. Early methods include black masking [16], blurring [31], and pixelation [13], while advanced approaches leverage inpainting techniques [6]. Inpainting can be diffusion-based [5], patch-based [12], or rely on machine learning models, such as CNNs [11] and GANs [28].

Despite advancements in privacy-preserving technologies, the development of scalable and interoperable solutions capable of adapting to the dynamic and evolving nature of the metaverse remains crucial [17], [32]. Specifically, to the best of our knowledge, there is no work experimenting combining different obfuscation and 3D reconstruction techniques for creating metaverse objects. In this context, our work aims to evaluate these combinations to provide practitioners with guidance on selecting the most effective strategy for generating privacy-preserving 3D scenes in the metaverse.

III. Study Design

To achieve the research objectives, we conducted a series of experiments to evaluate the impact of different obfuscation and 3D reconstruction techniques on the generation of privacy-compliant 3D models. We followed a structured process, systematically applying obfuscation techniques to remove privacy-sensitive elements from input images and reconstructing 3D models using various approaches. The different combinations of obfuscation and reconstruction techniques represented the independent variables, while a set of quality metrics computed on the final 3D models served as the dependent variables. By comparing these metrics across all approaches, we identified different technique combinations for various scenarios, providing valuable insights for both researchers and practitioners.

We structured the study as follows:

1) First, we recorded videos of objects containing privacy-threatening elements.
2) Second, we processed the recorded frames and applied various AI-based obfuscation approaches to remove or anonymize privacy-threatening elements.
3) Third, we used the obfuscated frames as input for different 3D reconstruction techniques to generate the corresponding 3D models. Additionally, to enable comparison, we also captured and reconstructed 3D models from the original object frames that did not contain privacy-sensitive objects.
4) Last, we computed a set of quality metrics for each reconstructed 3D model to determine whether significant differences existed between different technique combinations and to identify the most effective approaches for privacy-preserving 3D content generation.

A. Obfuscation Methods

We evaluated one algorithm for each of the four techniques introduced in Section II:

1) Black masking: Masking is performed by filling the private region with black pixels.
2) Blurring: Median blurring with a 25x25 kernel is implemented to better obscure the features of the hidden private object, unlike other blurring methods that may preserve sharp edges, potentially making the object recognizable.
3) Pixelation: Pixelation is performed by first resizing the frame to a lower resolution (64x64) and then scaling it back up using nearest-neighbor interpolation.
4) Navier Stokes Inpainting [5]: This algorithm is a traditional diffusion-based inpainting method that propagates edge and color information from the known region to the hidden region. It is implemented using OpenCV library.
5) Large Mask (LaMa) Inpainting [28]: This model is a machine learning-based inpainting model. LaMa leverages Fourier Convolutions for a wider receptive field, consequently generating contextually meaningful inpainting.

TABLE I
CATEGORIES OF PRIVACY-THREATENING ITEMS.

ID	Category	Example
C1	Nudity/Sexual	Photo of shirtless man or a couple
C2	Other people	Family photo
C3	Unorganized home	Trash or dirty desk
C4	Violence	Knife or gun
C5	Medical	Medicine or pill bottle
C6	Drinking/Party	Alcoholic drink or pack of cigarettes
C7	Appearance	Photo revealing a tattoo or messy hair
C8	Bad Character	Drugs or mugshot
C9	Religion/Culture	Religious book or crucifix
C10	Personal information	Credit card, ID, or passport

TABLE II
OBJECTS WITH PRIVACY-THREATENING ITEMS.

ID	Object	Privacy-Threatening items
O1	Cabinet	Knife (C4), alcoholic drink (C6), credit card (C10)
O2	Computer	Photo of a shirtless man with tattoos (C1. C7), Picture of a mugshot (C7, C8)
O3	Desk	Used paper napkin (C3), ID card (C2, C10), medicine, personal photo (C5), religious book (C9)

B. 3D Reconstruction Methods

The second step of the research process involved generating 3D models using different reconstruction techniques.

As discussed in Section II, the most commonly used techniques can be categorized into three major families: photogrammetry, 3DGS and NeRF-based methods. Since these approaches are well-established and extensively documented, we relied on past literature [27], [33], [35] to identify the most suitable implementations for our research. For both photogrammetry and 3DGS, we selected the Polycam tool [25]. Polycam was chosen for its efficiency and open-accessibility, offering a strong balance between accuracy and processing speed compared to its competitors [8]. Additionally, it supports 3D reconstruction from images or videos using either a photogrammetry or a 3DGS model, making it particularly well-suited to our needs. For NeRF-based reconstruction, we used Nerfacto [18], a hybrid method from the Nerfstudio Framework that combines NeRF with instant neural graphics primitives. It has proven to be among the top-performing NeRF implementations, enabling real-time use with high-quality results [27].

C. Experiment Procedure

Here we describe in detail the experiment procedure and justify the taken choices.

1) Step 1—Video Recording: We recorded videos of objects in a computer science research laboratory; we proceeded in this way for for two reasons: (1) to recreate a familiar yet complex environment, aligning with Dionisio's view of realism as key to the metaverse [10]; and (2) to leverage full control over the space, enabling scene complexity adjustments.

Regarding the selection of privacy-threatening items, the subjective nature of privacy has led to various studies aimed at identifying standard categories of privacy-sensitive objects [15], [36]. We adopted the classification proposed in [36], which is particularly suitable for our context, as it focuses on static 3D scenes containing inanimate private objects. Table I reports the ten identified categories, ranging from personal information (e.g., credit cards, IDs) to sensitive personal attributes (e.g., religious symbols, medical items).

Before proceeding with the main experiment, we conducted a feasibility pilot to ensure the correctness of the setup and to assess the required materials. Afterward, two researchers arranged the objects and placed privacy-threatening items on them according to the combinations listed in Table II. Once the setup was complete, we recorded videos by moving around the objects to capture all angles. Each object was filmed in two conditions: (1) without privacy-threatening items and (2) with the privacy-threatening items in place.

2) Step 2—Obfuscation and Frame Extraction: We processed the collected videos by applying different obfuscation techniques to remove privacy-threatening elements. We first applied segmentation techniques to detect sensitive elements and then used the obfuscation algorithms described in Section III-A to generate multiple versions of the videos, with each version corresponding to a different obfuscation method. The videos were subsequently converted into frame sequences for further processing.

3) Step 3—3D Reconstruction: We reconstructed 3D models of the objects from the frame sequences using the techniques described in Section III-B. Each reconstruction algorithm was configured at the highest level of detail to ensure that any reconstruction failures were not due to excessively low mesh resolution settings. Consequently, each model resulted from combining an obfuscation method and a 3D reconstruction method. Additionally, a "clean" model was generated for each object using a video that did not contain sensitive items, allowing for direct comparison. All the material produced from our experiments is available in our online appendix [24].

D. Data Analysis

In terms of metrics to assess model quality, based on the literature, we identified the following:

- **Usability—Number of Polygons**: We evaluated the usability of the models by analyzing the number of polygons that form the meshes generated by each model. Such a measure is crucial as it directly impacts real-time performances in metaverse environments [10]. Each polygon increases computational demands for rendering and lighting, requiring more hardware and software resources to ensure smooth simulation [7].
- **Realism—Structural Similarity Index (SSIM)**: We measured the realism of the reconstructed objects using SSIM, which quantifies perceptual similarity based on luminance, contrast, and structural integrity [34]. This metric allowed us to evaluate how closely the synthetic models resembled their real-world counterparts in terms of visual fidelity. We computed this metric by overlapping the picture of the real object with the rendering of the

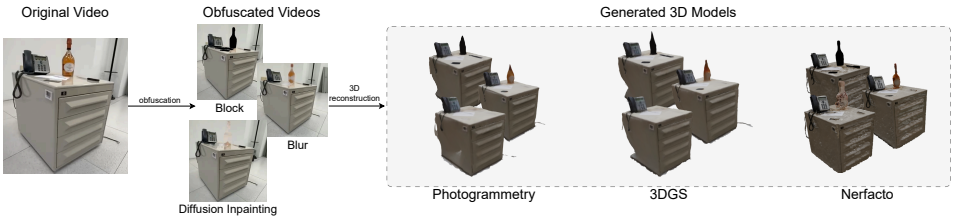

Fig. 1. Examples of the generation pipeline for the cabinet object.

corresponding 3D-reconstructed model. A higher SSIM value indicates greater similarity between images.

- **Privacy Level—Hausdorff Distance**: The Hausdorff distance [3] is a metric for 3D cloud comparison. We used it to measure the similarity between only the private objects in each reconstructed model and the original model—i.e., scanned with sensitive elements. For example, for the cabinet object, we computed the Hausdorff distance between the knife, alcoholic drink, and credit card objects before they were obfuscated (as they were scanned initially), and after they were obfuscated. A lower distance indicates greater similarity between the models. The combination of the methods that ensures the highest level of privacy is the one that distorts the private objects the most, making them as different as possible from how they initially looked.

These metrics were computed for each model generated through the combination of obfuscation and reconstruction techniques and subsequently compared.

IV. Analysis of the Results

The experiment produced a series of 3D models reconstructed from real objects in the scanned office scene as described in Section III-C3. Figure 1 illustrates examples of the followed pipeline for the cabinet object; Table III and IV summarize the obtained results.

*a) **Number of Polygons**:* Meshes were analyzed by polygon count to assess real-time rendering suitability. The results indicate that no combination led to significantly fewer polygons. In particular, Nerfacto consistently produced detailed models and accurately reconstructed flat, uniform surfaces, something photogrammetry and 3DGS failed to achieve, especially for the desk object. While 3DGS slightly outperformed photogrammetry in polygon reduction for the cabinet object (but not for the computer object), the difference was minor. Nerfacto had the highest polygon count, yielding $\approx$4.7–6.5$\times$ and $\approx 1.2 - 1.7\times$ more polygons than the lighter Photogrammetry/3DGS versions for the cabinet and computer objects, respectively, but captured object shapes more completely. These results highlight two key points: (1) photogrammetry and 3DGS yield lighter models but struggle with smooth surfaces; (2) Nerfacto better captures shapes, including plain ones, at the cost of higher polygon density.

*b) **SSIM**:* We employed the SSIM metric to evaluate the realism of the objects. The measurements confirmed that no clear significant pattern emerges in different obfuscation and reconstruction techniques; rather, differences are observed across objects. Specifically, the cabinet object exhibits the lowest SSIM values (< 0.21), as its model undergoes a greater mesh deformation due to its rectangular shape and glossy, light-colored material. In contrast, the computer object, which is darker and contains more details, substantially increases the SSIM metric ($\approx$0.50–0.56). The SSIM values obtained are even higher for the desk object reconstructed by Nerfacto ($\approx$0.75-0.9), suggesting a high level of realism achieved for such object. This high variability of SSIM ($\approx 2.4 - 3.5\times$ difference between objects) suggests that the primary factor influencing performance is more closely related to the shape of the object under examination.

*c) **Hausdorff Distance**:* The level of privacy was assessed using the Hausdorff Distance between the mesh of the sensitive objects reconstructed from the non-obfuscated models and those generated from obfuscated references using various methods. For the cabinet object, Diffusion Inpainting achieved the highest distances ($\approx$166-177% increase), particularly when paired with 3DGS and Photogrammetry. Conversely, Nerfacto consistently produced lower Hausdorff values across cabinet ($\approx$50-475% decrease) and desk scenarios, suggesting reduced privacy performance on volumetric or spatially complex objects. This could be due to its higher capability of capturing objects' details even when they are obfuscated, consequently resulting in a lower privacy preservation. For the computer object, of which sensitive items primarily consisted of flat shapes, the Hausdorff values yielded more similar results, with minimal differences (< 0.05) among the values. Finally, for the desk object—for which only Nerfacto models were available—Pixelate and Diffusion Inpainting emerged as the most effective. These findings suggest that while Nerfacto can retrieve more accurate and detailed 3D reconstructions, even where Photogrammetry and 3DGS fail, it ultimately results in a less-preserving privacy method, with Photogrammetry and 3DGS offering higher distances from the sensitive objects and stronger privacy preservation.

V. Discussion

Our preliminary experimental results pave the way for various discussion points. Figure 2 shows visual examples of the obtained models, which we will discuss in the following.

Reconstruction techniques: shape preservation vs. object complexity. Nerfacto generally outperformed other reconstruction methods in preserving object geometry, although at the cost of a higher polygon count. As shown in Block

TABLE III
COMPARISON OF NUMBER OF POLYGONS, SSIM, AND HAUSDORFF DISTANCE FOR DIFFERENT RENDERING METHODS.

Object	Obfuscation	Number of Polygons			SSIM			Hausdorff Distance		
		Photogr.	3DGS	Nerfacto	Photogr.	3DGS	Nerfacto	Photogr.	3DGS	Nerfacto
Cabinet	**Pixelate**	8,941	8,686	56,614	0.154	0.154	0.140	0.167	0.043	0.029
	Blur	8,690	8,315	47,871	0.163	0.153	0.154	0.044	0.047	0.048
	Block	8,176	**7,904**	38,431	0.138	0.164	0.129	0.035	0.055	0.037
	Diffusion Inpaint	8,790	8,038	41,549	0.160	0.151	0.164	0.173	**0.180**	0.065
	ML-based Inpaint	8,551	8,334	8,059	0.194	**0.208**	0.179	0.169	0.168	0.060
Computer	**Pixelate**	13,066	20,110	15,944	0.534	0.540	0.517	0.024	0.028	0.030
	Blur	12,331	13,753	16,432	0.537	0.525	0.498	**0.054**	0.015	0.031
	Block	12,770	**11,425**	17,345	0.520	0.524	0.507	0.050	0.042	0.031
	Diffusion Inpaint	12,859	13,462	21,913	**0.560**	0.523	0.520	0.015	0.016	0.037
	ML-based Inpaint	15,376	15,980	21,101	0.550	0.528	0.539	0.025	0.013	0.049

Fig. 2. Example of results. Block A compares the reconstruction of the Desk in a successful case using Nerfacto and a failed reconstruction using photogrammetry. Block B compares three reconstructions of the Cabinet with the pixelate obfuscation technique. Block C compares three obfuscation methods applied to a flat image of the Computer screen.

TABLE IV
DESK OBJECT'S METRICS FOR THE NERFACTO METHOD.

Obfuscation	N. of Polygons	SSIM	Hausdorff Dist.
Pixelate	17,421	0.761	**0.049**
Blur	**14,569**	0.759	0.034
Block	15,495	0.765	0.034
Diffusion Inpaint	16,658	0.762	**0.049**
ML-based Inpaint	10,012	**0.799**	0.034

A, Nerfacto is the only method able to reconstruct the desk object successfully, whereas photogrammetry completely failed due to the flat and glossy surface. Nevertheless, the success of each technique depends on the shape and texture properties of the object. Detailed and textured items—such as the computer—led to better results and higher SSIM values. Conversely, objects with smooth, reflective surfaces—like the desk or cabinet—posed significant challenges, resulting in incomplete reconstructions or visible mesh artifacts. No single technique consistently outperformed others across all objects, underscoring the absence of a universally optimal obfuscation-reconstruction combination. Instead, certain methods proved more effective for specific object types (e.g., black block for computers, blurring for flat surfaces). These findings point to the need for adaptive or ensemble-based approaches that tailor the processing pipeline to the characteristics of each object.

Obfuscation techniques and their interaction with reconstruction. The effectiveness of each obfuscation technique is not only object-dependent but also influenced by the reconstruction method. Block B illustrates this point: the same obfuscation method (pixelate) produces entirely different outcomes across reconstructions. In the photogrammetry model, the bottle is distorted but recognizable; in Nerfacto, the object is nearly reconstructed in full; while in 3DGS, it is completely missing, indicating a higher privacy level. Similarly, Block C shows the effect of obfuscation techniques on flat images. Pixelate retains facial contours and texture details, making the subject still identifiable; block masking hides facial information while preserving the silhouette; diffusion inpainting results in an intermediate outcome, blending color cues with partial removal of the original shape. These results suggest that the success of an obfuscation method cannot be evaluated in isolation—it must be assessed in the context of the target reconstruction strategy.

Implications for privacy-aware metaverse generation. The findings reinforce the importance of combining obfuscation and reconstruction techniques thoughtfully in metaverse applications, where realism and privacy must coexist. The ability to automatically remove sensitive elements from input captures while maintaining plausible 3D reconstructions is crucial for non-expert users creating virtual spaces. These preliminary results mark a step forward toward the development of virtual environments that incorporate real-world captures, allowing users to replicate spaces while safeguarding security and privacy digitally.

VI. CONCLUSION AND FUTURE WORK

Our findings reveal that no single combination of techniques consistently outperforms others in all cases for the task of creating privacy-preserving 3D models, of good quality and usability, for the metaverse. Instead, the effectiveness of each method depends on the object's characteristics and the specific setting, demonstrating that different obfuscation-reconstruction strategies may be better suited to different scenarios.

Future research can build upon these results by exploring the correlation between object properties and the performance of obfuscation and reconstruction techniques to select the most appropriate techniques based on the specific context. AI-driven recommendation mechanisms could be incorporated to predict and select the optimal obfuscation–reconstruction combination based on the given scene and object features.

Moreover, addressing internal validity concerns, such as bias introduced by manual placement of privacy-sensitive items, will require automating or randomizing scene setup to improve reproducibility and minimize bias. Strengthening external validity will involve evaluating the pipeline across diverse real-world environments, ranging from home offices to public and outdoor spaces, to capture the variability of privacy threats in different contexts. Finally, enhancing construct validity will necessitate expanding beyond current metrics by developing more objective, standardized measures of privacy and realism, and by conducting comprehensive user studies to align technical evaluations with human perceptions.

VII. Acknowledgement

We thank the AUB student Mohamad Daouk for his help in segmenting the videos.

References

[1] M. Aharchi and M. Ait Kbir. A review on 3d reconstruction techniques from 2d images. In *Innovations in Smart Cities Applications Edition 3: The Proceedings of the 4th International Conference on Smart City Applications 4*, pages 510–522. Springer, 2020.

[2] T. Amano, T. Mizumoto, S. M. Kala, H. Yamaguchi, T. Matsui, and K. Yasumoto. Visual privacy control for metaverse and the beyond. *IEEE Pervasive Computing*, 23:10–17, 2024.

[3] N. Aspert, D. Santa-Cruz, and T. Ebrahimi. Mesh: Measuring errors between surfaces using the hausdorff distance. In *Proceedings. IEEE international conference on multimedia and expo*, volume 1, pages 705–708. IEEE, 2002.

[4] J. Baqersad, P. Poozesh, C. Niezrecki, and P. Avitabile. Photogrammetry and optical methods in structural dynamics–a review. *Mechanical Systems and Signal Processing*, 86:17–34, 2017.

[5] M. Bertalmio, A. Bertozzi, and G. Sapiro. Navier-stokes, fluid dynamics, and image and video inpainting. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 1, pages I–I, 2001.

[6] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, page 417–424, USA, 2000. ACM Press/Addison-Wesley Publishing Co.

[7] M. Botsch, L. Kobbelt, M. Pauly, P. Alliez, and B. Lévy. *Polygon mesh processing*. CRC press, 2010.

[8] M. M. Buzayan, A. H. Elkezza, S. F. Ahmad, N. M. Salleh, and I. Sivakumar. A comparative evaluation of photogrammetry software programs and conventional impression techniques for the fabrication of nasal maxillofacial prostheses. *The Journal of Prosthetic Dentistry*, 2023.

[9] J. D. N. Dionisio, W. G. B. Iii, and R. Gilbert. 3d virtual worlds and the metaverse: Current status and future possibilities. *ACM Computing Surveys (CSUR)*, 45(3):1–38, 2013.

[10] J. D. N. Dionisio, W. G. B. Iii, and R. Gilbert. 3d virtual worlds and the metaverse: Current status and future possibilities. *ACM Computing Surveys (CSUR)*, 45(3):1–38, 2013.

[11] Z. Guo, Z. Chen, T. Yu, J. Chen, and S. Liu. Progressive image inpainting with full-resolution residual network. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 2496–2504. ACM, 2019.

[12] J. Herling and W. Broll. High-quality real-time video inpaintingwith pixmix. *IEEE Transactions on Visualization and Computer Graphics*, 20(6):866–879, 2014.

[13] S. Hill, Z. Zhou, L. Saul, and H. Shacham. On the (in) effectiveness of mosaicing and blurring as tools for document redaction. *Proceedings on Privacy Enhancing Technologies*, 2016.

[14] J. Jam, C. Kendrick, K. Walker, V. Drouard, J. G.-S. Hsu, and M. H. Yap. A comprehensive review of past and present image inpainting methods. *Computer Vision and Image Understanding*, 203:103147, 2021.

[15] Y. Li, N. Vishwamitra, H. Hu, and K. Caine. Towards a taxonomy of content sensitivity and sharing preferences for photos. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–14, New York, NY, USA, 2020. Association for Computing Machinery.

[16] Y. Li, N. Vishwamitra, B. P. Knijnenburg, H. Hu, and K. Caine. Blur vs. block: Investigating the effectiveness of privacy-enhancing obfuscation for images. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1343–1351, 2017.

[17] Y. Lu, B. Deng, Z. Zhong, T. Zhang, Y. Quan, H. Cai, and S. He. 3d snapshot: Invertible embedding of 3d neural representations in a single image. *IEEE transactions on pattern analysis and machine intelligence*, PP, 2024.

[18] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.

[19] S. Mystakidis. Metaverse. *Encyclopedia*, 2(1):486–497, 2022.

[20] Y. Otoum, N. Gottimukkala, N. Kumar, and A. Nayak. Machine learning in metaverse security: Current solutions and future challenges. *ACM Computing Surveys*, 56:1 – 36, 2024.

[21] J. R. Padilla-López, A. A. Chaaraoui, and F. Flórez-Revuelta. Visual privacy protection methods: A survey. *Expert Systems with Applications*, 42(9):4177–4195, 2015.

[22] V. Paneva, M. Strauss, V. Winterhalter, S. Schneegass, F. Alt, and F. Alt. Privacy in the metaverse:. *IEEE Pervasive Computing*, 23:73–78, 2024.

[23] V. Pentangelo, D. Di Dario, V. De Martino, M. D. Buono, and S. Lambiase. Accelerating 3d scene development for the metaverse: Lessons from photogrammetry and manual modeling. In *2024 2nd International Conference on Intelligent Metaverse Technologies & Applications (iMETA)*, pages 190–197. IEEE, 2024.

[24] V. Pentangelo, S. Tabet, A. Kayssi, I. H. Elhajj, and F. Palomba. Privacy-aware 3d reconstruction and obfuscation in the metaverse. https://figshare.com/s/e13c9c9812955be9b639.

[25] Polycam. Polycam. https://poly.cam [Accessed: May 2024].

[26] F. Remondino, A. Karami, Z. Yan, G. Mazzacca, S. Rigon, and R. Qin. A critical analysis of nerf-based 3d reconstruction. *Remote Sensing*, 15(14):3585, 2023.

[27] F. Remondino, A. Karami, Z. Yan, G. Mazzacca, S. Rigon, and R. Qin. A critical analysis of nerf-based 3d reconstruction. *Remote Sensing*, 15(14):3585, 2023.

[28] R. Suvorov, E. Logacheva, A. Mashikhin, A. Remizova, A. Ashukha, A. Silvestrov, N. Kong, H. Goka, K. Park, and V. Lempitsky. Resolution-robust large mask inpainting with fourier convolutions, 2021.

[29] S. Tabet, A. Kayssi, and I. H. Elhajj. Ai-enhanced mobile diminished reality for preserving 3d visual privacy. In *2024 2nd International Conference on Intelligent Metaverse Technologies & Applications (iMETA)*, pages 141–148. IEEE, 2024.

[30] J. Tang, K. Fan, W. Yin, S. Yang, Y. Huang, A. Liu, N. N. Xiong, M. Dong, T. Wang, and S. Zhang. A quality-aware and obfuscation-based data collection scheme for cyber-physical metaverse systems. *ACM Transactions on Multimedia Computing, Communications and Applications*, 2024.

[31] J. Tekli, B. al Bouna, R. Couturier, G. Tekli, Z. al Zein, and M. Kamradt. A framework for evaluating image obfuscation under deep learning-assisted privacy attacks. In *2019 17th International Conference on Privacy, Security and Trust (PST)*, pages 1–10, 2019.

[32] C. Y. Wang, S. Sriram, and A. S. Won. Shared realities: Avatar identification and privacy concerns in reconstructed experiences. *Proceedings of the ACM on Human-Computer Interaction*, 5:1 – 25, 2021.

[33] X. Wang, C. Wang, B. Liu, X. Zhou, L. Zhang, J. Zheng, and X. Bai. Multi-view stereo in the deep learning era: A comprehensive review. *Displays*, 70:102102, 2021.

[34] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

[35] T. Wu, Y.-J. Yuan, L.-X. Zhang, J. Yang, Y.-P. Cao, L.-Q. Yan, and L. Gao. Recent advances in 3d gaussian splatting. *Computational Visual Media*, 10(4):613–642, 2024.

[36] C. Zhao, J. Mangat, S. Koujalgi, A. Squicciarini, and C. Caragea. Privacyalert: A dataset for image privacy prediction. *Proceedings of the International AAAI Conference on Web and Social Media*, 16(1):1352–1361, May 2022.