Visualization of 3D Ultrasound Uterine Data in Virtual Reality

Ilija Gavrilović[†], Ciril Bohak[*,†]

University of Ljubljana, Faculty of Computer and Information Science, Večna pot 113, SI-1000 Ljubljana, Slovenia

Abstract

Virtual reality (VR) has emerged as a powerful medium for medical imaging visualization and training, offering immersive and interactive experiences that complement conventional two-dimensional (2D) imaging workflows. In gynecology, ultrasound (US) is the primary imaging modality for assessing uterine morphology, yet its interpretation requires significant spatial reasoning and is challenging for trainees with limited clinical exposure. This paper presents a standalone VR system for interactive visualization of uterine US volumes, designed for deployment on Meta Quest 2. The system integrates three core components: (i) textured surface meshes mapped directly from US intensity values, (ii) per-vertex deviation heatmaps comparing individual anatomy to a population-average uterus, and (iii) orthogonal slice browsing with adjustable transfer functions and lookup tables. The system is developed in Unity for immersive rendering. Preliminary demonstrations indicate that providing outside- and inside-based views improves spatial understanding and provides educational value by contextualizing individual variability. Performance profiling confirms real-time rendering on standalone hardware, ensuring fluid interaction without tethered computing resources. By unifying segmentation-driven shape analysis with immersive visualization, this work highlights the potential of lightweight VR applications to enhance gynecological training and provide accessible platforms for medical education and research.

Keywords

virtual reality, histeroscopy, simulation, uterus.

1. Introduction

Virtual Reality (VR) has matured into a practical medium for interactive visualization and simulation across various domains, including medical education and training. In gynecology, many diagnostic and therapeutic procedures are invasive, opportunities for repeated hands-on practice with real patients are limited, and the interpretation of two-dimensional (2D) imaging modalities such as Ultrasound (US) requires substantial spatial reasoning. These constraints motivate the development of realistic, repeatable, and safe training environments where learners can explore anatomy, rehearse procedures, and build visuospatial expertise without risk to patients. VR-based simulation meets these needs by enabling controlled scenarios, objective progress tracking, and immersive interaction with complex anatomical data [1, 2, 3].

Hysteroscopy—a minimally invasive endoscopic procedure for examining the uterine cavity—is a representative task that particularly benefits from immersive simulation. Successful navigation of the cervical canal and uterine cavity requires an accurate understanding of individual uterine morphology and spatial orientation, both of which are difficult to master through conventional 2D ultrasound training alone. Conventional training approaches rely heavily on observation and limited hands-on opportunities, which can result in uneven skill acquisition and reduced learner engagement. An immersive visualization environment that fuses volumetric US data with a patient-specific surface representation can help bridge the gap between 2D image slices and three-dimensional (3D) anatomical understanding, offering a cognitively intuitive way to link imaging with real procedural perspectives.

Human-Computer Interaction Slovenia 2025, October 13, 2025, Koper, Slovenia

[*]Corresponding author.
[†]These authors contributed equally.
ig80333@student.uni-lj.si (I. Gavrilović); ciril.bohak@fri.uni-lj.si (C. Bohak)
https://lgm.fri.uni-lj.si/ciril (C. Bohak)
0000-0002-9015-2897 (C. Bohak)

https://doi.org/10.26493/978-961-293-559-7.5

A key educational and clinical challenge is understanding the variability in uterine morphology across patients. Differences in uterine size, shape, and orientation can influence both diagnostic interpretation and the execution of procedures such as hysteroscopy or intrauterine device (IUD) placement. Recognizing and comparing these variations supports the development of spatial reasoning, facilitates individualized treatment planning, and enhances awareness of normal versus pathological anatomical configurations. Therefore, a visualization system that allows comparison between an individual uterus and a population-based morphological reference can serve as a valuable tool for both training and research.

Building on these motivations, this work presents a VR application for interactive visualization of uterine US volumes that integrates: (i) a surface mesh view textured with information derived directly from the US volume, (ii) per-vertex visualization of distances to a population-average uterine shape, and (iii) slice-based exploration of the volume along orthogonal planes, enabling users to cross-reference surface structures with internal volumetric information. The inclusion of slice-based exploration is particularly important for training, as it reinforces the connection between volumetric US data and the spatial geometry visualized on the uterine surface. The system is designed for standalone deployment on Meta Quest 2, emphasizing fluid interaction, spatial understanding, and real-time performance suitable for instructional and clinical educational contexts.

2. Related Work

Extended Reality (XR) technologies, encompassing VR, Augmented Reality (AR), and Mixed Reality (MR), have become increasingly prominent in medical education and clinical visualization. Prior work has established the potential of XR in medical visualization [1, 2, 4]. Immersive environments offer unique advantages for exploring complex anatomical structures, improving spatial understanding, and supporting procedural training. In particular, VR enables embodied interaction with volumetric datasets, allowing clinicians and trainees to intuitively navigate anatomy beyond the limitations of flat, screen-based tools. Recent research emphasizes the importance of integrating multiple data representations—specifically, linking slice-based 2D information with immersive 3D views—to reinforce spatial reasoning and diagnostic accuracy [3]. Combining orthogonal plane navigation with volumetric exploration allows learners to understand the correspondence between image slices and three-dimensional form, a critical skill in interpreting ultrasound data.

The preparation of medical data for immersive visualization remains a significant technical challenge. Automatic segmentation of anatomical structures from ultrasound volumes is particularly demanding due to speckle noise, acoustic shadowing, and low tissue contrast. Deep learning architectures such as U-Net [5] and nnU-Net [6] have become the standard for biomedical image segmentation, while hybrid and attention-based variants such as RA-UNet [7] and generalization-focused approaches [8] address domain shifts and noisy acquisitions. Survey work highlights the persistence of imperfect datasets and the importance of developing robust segmentation frameworks for clinical deployment [9]. In gynecology, segmentation and alignment techniques have recently been applied to ultrasound volumes of the uterus, enabling the construction of population-level statistical models that capture inter-patient morphological variability [10]. The open-source dataset *UterUS* provides both data and implementations for these workflows, promoting reproducibility and supporting downstream applications in visualization and training.

Several studies have explored the use of VR in gynecological education and simulation. Early efforts demonstrated the feasibility of web-based VR environments for medical visualization [11], while later systems extended toward interactive navigation of volumetric and patient-specific data [12]. In the context of hysteroscopy, immersive simulation is particularly valuable, as clinicians must reason spatially within a confined anatomical space. Recent work on automatic uterine segmentation and geometric alignment [10] directly supports this domain by providing accurate surface and volumetric models suitable for subject-specific visualization and population-level comparison.

The integration of artificial intelligence into ultrasound imaging pipelines further enhances data preparation for immersive visualization [13]. AI-driven segmentation facilitates the generation of accurate anatomical models, while XR visualization offers an intuitive environment for validating and interpreting algorithmic results. This synergy enables not only the inspection of individual anatomy but also the exploration of population-level variability—providing a pedagogical bridge between data analysis and experiential learning.

In summary, prior work has established the educational and clinical potential of XR visualization, demonstrated robust AI-based segmentation pipelines, and explored the foundations of population-based uterine modeling. However, few systems integrate segmentation-driven shape analysis with interactive, immersive visualization on standalone VR hardware. The present work addresses this gap by combining surface- and volume-based representations with population-level deviation mapping, offering an accessible and pedagogically oriented visualization framework for gynecological training.

3. Methodology

Our methodology follows a modular pipeline that transforms raw US volumes into an interactive VR environment optimized for standalone deployment. The workflow comprises three main components: (i) data preparation, (ii) surface and texture processing, and (iii) immersive visualization and interaction design. Each stage builds on open-source frameworks to ensure reproducibility and extendability.

3.1. Data Preparation

Raw 3D US volumes of the uterus serve as the input data. Segmentation and alignment are performed using the publicly available *UterUS* pipeline [10], which provides preprocessed meshes and population-based shape correspondence. This stage yields polygonal surface representations of individual uteri registered to a population-average model, enabling consistent visualization of morphological variability across subjects. Because the segmentation framework and dataset are described in prior work, this paper focuses on the visualization and interaction stages that build upon these outputs.

3.2. Processing the Surface Mesh

Segmented uterine meshes are refined for real-time rendering. Each mesh is registered to the population-average uterus, and per-vertex deviations are encoded as scalar values that later serve as color-coded indicators of local anatomical differences. In Blender, UV unwrapping and texture baking are performed to map ultrasound-derived intensity values onto the surface. Surface normals are recalculated to ensure coherent lighting for both external and internal perspectives. The processed assets—including meshes, deviation fields, and textures—are exported in optimized formats for Unity to ensure efficient rendering on standalone VR hardware.

3.3. Volumetric Data Integration

To maintain clinical relevance and aid spatial reasoning, orthogonal ultrasound slice views (axial, sagittal, and coronal) are integrated alongside the surface model. Slice stacks are aligned with the registered meshes and stored as 2D textures for interactive browsing within the VR environment. Users can scroll through slices using controller input while simultaneously inspecting the corresponding anatomical regions on the surface model. A transfer-function module enables dynamic intensity and transparency adjustments; detailed visualization parameter settings are discussed in the Results section.

3.4. Immersive Visualization in Unity

The immersive environment was implemented in Unity, leveraging its built-in XR toolkit and native support for the Meta Quest 2 headset. Two complementary visualization modes are provided: (i) an **external view** for inspecting overall uterine morphology and population deviation mapping, and (ii) an **internal (hysteroscopic) view**, where the virtual camera is positioned inside the uterine cavity to simulate endoscopic navigation.

Interaction within the VR environment combines ray-based User Interface (UI) controls with direct 3D manipulation. Users can grab, rotate, and scale the uterine model using handheld controllers, switch seamlessly between external and internal viewpoints, and move through the cavity using joystick-based locomotion or teleportation. Slice planes can be activated, repositioned, or scrolled interactively to correlate volumetric US slices with the surface model, supporting the development of spatial reasoning between 2D and 3D representations. Menu-based controls allow users to adjust visualization options such as color mapping, transparency, and population deviation overlays.

To maintain real-time performance on standalone hardware, mesh decimation and texture compression were applied to balance geometric fidelity and memory footprint. The final build sustains 72 FPS on Meta Quest 2, providing a stable and low-latency experience suitable for educational use.

A preliminary user evaluation was conducted with eight participants: four gynecology residents, two medical imaging researchers, and two computer science students with prior VR experience. Participants were given a short tutorial on navigation and were then asked to complete three tasks: (1) identify major anatomical landmarks in the external view, (2) navigate from the cervical canal to the fundus in the internal view, and (3) locate corresponding structures in the slice-based visualization. After completing the tasks, participants answered four open-ended questions focusing on usability, perceived educational value, clarity of anatomical representation, and comfort during use. No small-scale questionnaires (e.g., Likert items) were used at this stage to encourage unconstrained qualitative feedback. To minimize experimenter bias, all instructions were standardized, and participants were free to explore the application for up to 10 minutes per session.

Feedback varied according to background. Medical participants emphasized the value of integrating 2D slice information with immersive 3D visualization for understanding uterine orientation and depth. In contrast, computer science participants commented on rendering quality, frame rate, and interface responsiveness. Across all groups, users appreciated the ability to freely switch between internal and external perspectives, noting that it improved comprehension of the spatial relationships between ultrasound slices and surface geometry.

These insights will inform the design of a structured usability study in future work, incorporating quantitative metrics and standardized questionnaires to evaluate learning outcomes and interaction efficiency.

4. Results

This section presents the qualitative and quantitative evaluation of the proposed VR system for interactive visualization of uterine US volumes. We focus on (i) visual fidelity and correctness of the rendered anatomy, (ii) performance of the standalone application on Meta Quest 2, and (iii) educational affordances observed during preliminary demonstrations.

4.1. Qualitative Visualization Outcomes

Figure 1 illustrates the external mesh visualization with textures derived from US intensity values. Surface deviations from the population-average uterus are encoded as a color-coded heatmap, enabling immediate identification of morphological differences. In the external view, clinicians can readily inspect the global shape and surface variability, while deviations larger than 3 mm are clearly visible as localized hot spots.

The internal mode (Fig. 2) provides a first-person perspective within the uterine cavity, mimicking hysteroscopic navigation. This view is particularly effective in demonstrating the geometry of the cavity and cervical canal. Dynamic lighting and correct surface normals ensure realistic shading and depth perception, supporting spatial understanding in an immersive context.

Orthogonal slice browsing is shown in Fig. 3. Users can scroll through axial, sagittal, and coronal planes in real time, directly linking volumetric intensities with anatomical landmarks on the surface

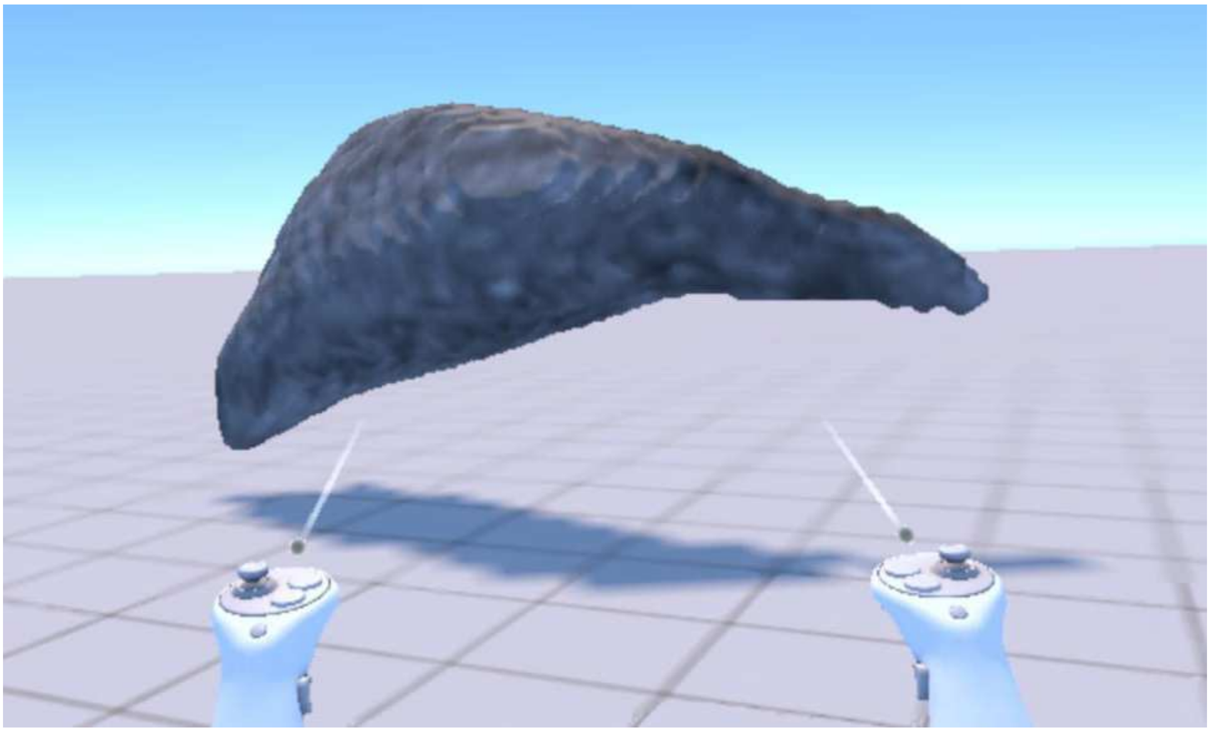

Figure 1: External view of the uterus surface mesh: texture-mapped ultrasound intensities.

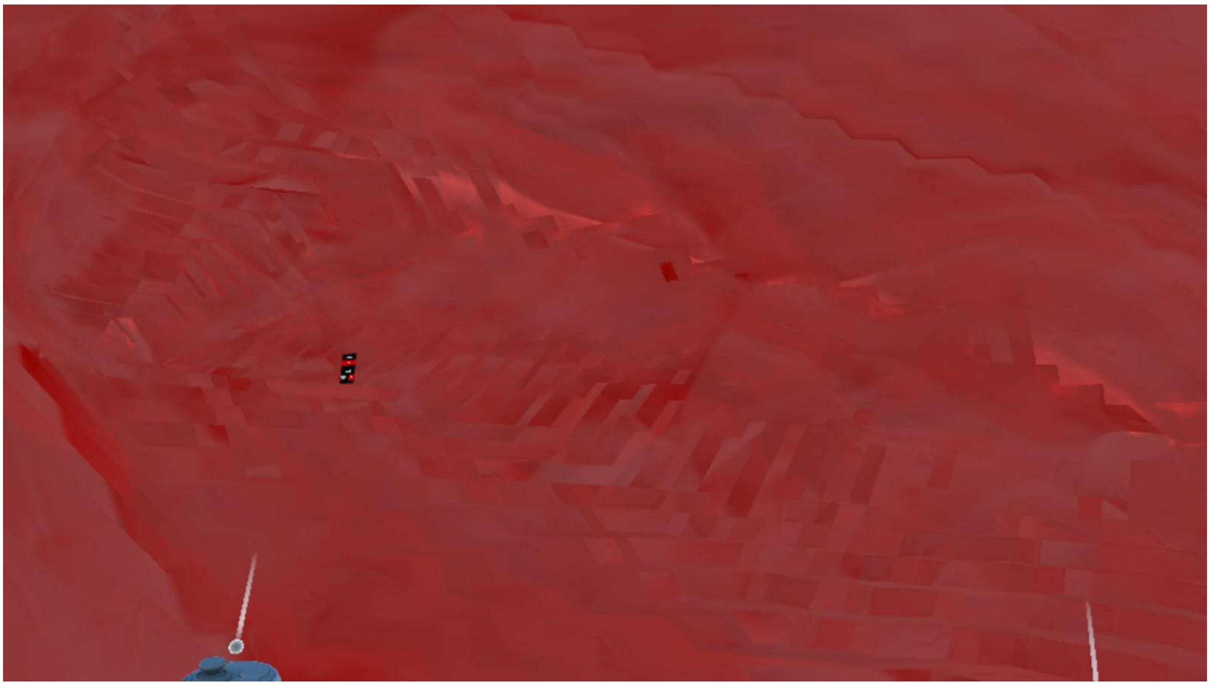

Figure 2: Internal view of the uterus cavity. Clinicians can explore the cervical canal and endometrial cavity in a manner analogous to hysteroscopic navigation.

mesh. Windowing, brightness, and gamma adjustments allow clinicians to adapt the visualization to individual preferences, reflecting common workflows in conventional ultrasound analysis.

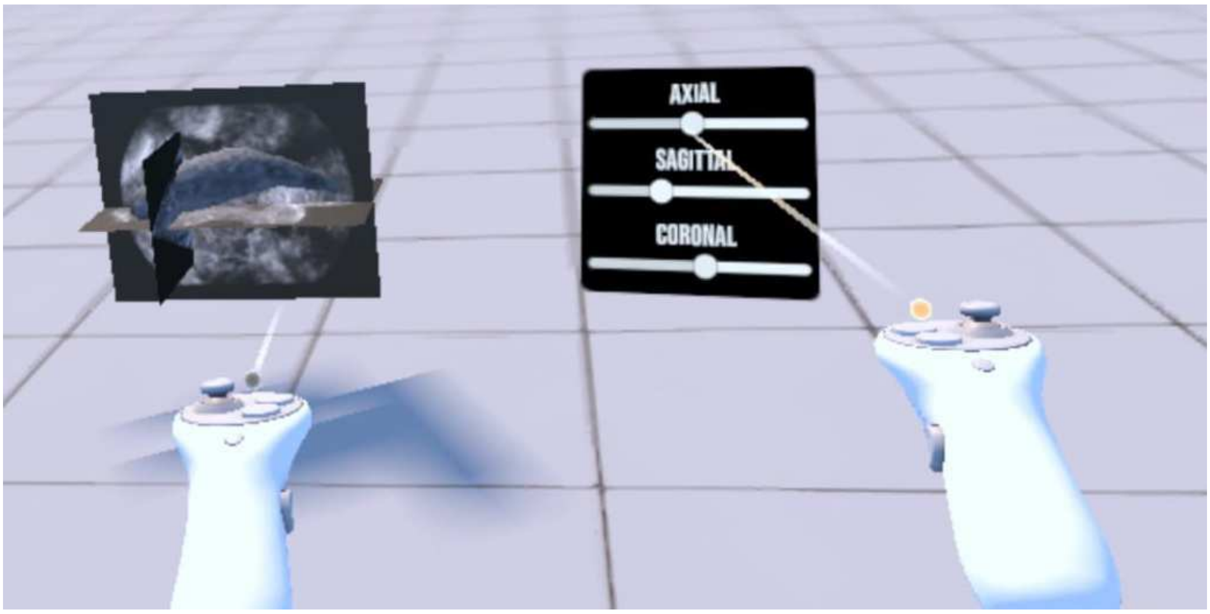

Figure 3: Orthogonal slice views (axial, sagittal, coronal) in VR.

4.2. Transfer Function Adjustment

An additional feature of the system is the real-time adjustment of the transfer function in the outside and inside view visualizations. Users can interactively modify window center, window width, brightness, and gamma values, which directly affect the contrast and visibility of structures within the US data. Custom look-up tables with alpha channels enable semi-transparent renderings, allowing subtle anatomical boundaries to be emphasized without occluding surrounding structures. The adjustments are illustrated in Fig. 4 for both inside and outside views.

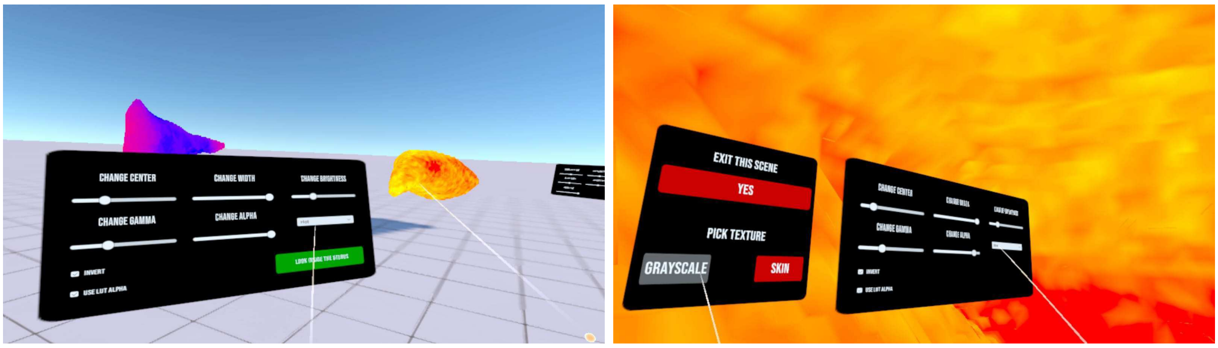

Figure 4: Adjustment of the transfer function for outside (left) and inside (right) views.

4.3. Performance on Meta Quest 2

Performance profiling was conducted using Unity's built-in frame diagnostics on Meta Quest 2. The application maintained an average of 71.8 FPS across all tested scenarios, with occasional dips below 70 FPS only when switching between external and internal viewpoints with high-resolution textures enabled. Latency remained below 20 ms, which is acceptable for extended VR sessions without inducing discomfort. Texture resolution (up to 2048^2) and mesh density (~40k vertices after decimation) proved sufficient to preserve anatomical detail while sustaining real-time performance. The application requested under 2 GB of RAM and ∼1.5 GB of VRAM. The usage stabilized throughout the use and did not increase with longer use.

4.4. Educational Affordances

Preliminary demonstrations with graduate students in biomedical engineering and computer science indicated that the combined slice and mesh visualization improved comprehension of uterine shape and variability. Participants reported that internal viewpoint navigation enhanced their spatial intuition about the cavity and cervical canal, which is difficult to achieve with monoscopic 2D images. The deviation heatmap was highlighted as a valuable feature for understanding anatomical variability across a population, suggesting potential for use in both medical training and research contexts.

5. Discussion

The presented system demonstrates how immersive visualization of uterine US data can be achieved on standalone VR hardware through a modular pipeline integrating segmentation, surface processing, and interactive rendering. This section discusses the educational implications, technical contributions, AI-assisted visualization components, and current limitations, followed by future research directions.

5.1. Educational Impact and Design Rationale

The primary motivation for this system is to address challenges in gynecological education, particularly the need for safe, repeatable, and spatially coherent training tools. By combining textured surface meshes, per-vertex deviation heatmaps, and interactive slice exploration, the application bridges conventional 2D ultrasound interpretation with immersive 3D anatomical understanding. This multimodal visualization enables trainees to directly correlate features visible in 2D slices with their spatial location on the uterus, reinforcing visuospatial reasoning—an essential skill for hysteroscopic procedures and diagnostic imaging. Early user feedback suggests that this dual representation improves engagement and helps learners contextualize anatomical variability in relation to population norms.

5.2. Technical and Hardware Contributions

A key strength of the approach is the use of standalone VR hardware, specifically Meta Quest 2, to deliver high-fidelity visualization without reliance on external computing resources. This portability lowers technical and logistical barriers to adoption in educational or clinical settings. Consistent real-time performance above 70 FPS ensures smooth, low-latency interaction, reducing motion sickness and maintaining user comfort during extended training sessions. Performance optimization through mesh decimation and texture compression confirms that complex medical visualization tasks can be efficiently executed on accessible consumer-grade devices.

5.3. AI-Integrated Visualization Advantages

The integration of segmentation-driven shape modeling provides an additional layer of pedagogical value. By leveraging population-level shape correspondence derived from automated segmentation pipelines [10], the system situates each patient-specific uterus within a broader statistical model. This allows learners to visually compare individual morphology with typical anatomical variations, fostering a deeper understanding of normal versus atypical cases. Unlike most XR-based visualization systems that focus solely on volumetric rendering [1, 2, 3], the presented framework explicitly links per-patient data to explainable, population-level information—aligning with current trends in interpretable medical AI [13].

5.4. Limitations

Despite promising results, several limitations remain. First, the accuracy of the visualization pipeline is constrained by segmentation quality. Although architectures such as U-Net [5] and nnU-Net [6] provide robust baselines, errors in automated segmentation may propagate to the resulting meshes, potentially misrepresenting fine anatomical details. Second, the current dataset is limited in size and diversity, which restricts the generalizability of findings. Third, user evaluation has been limited to preliminary qualitative feedback rather than structured usability or clinical validation studies. Finally, while the system achieves real-time performance on standalone hardware, future iterations with higher-resolution models or dynamic segmentation may challenge device capacity. Future work will address these limitations through expanded datasets, improved preprocessing, and larger-scale evaluations.

5.5. Future Work

Planned extensions will focus on both educational validation and technical enhancement. First, structured clinical studies involving gynecologists and medical students will assess usability, engagement, and learning outcomes using standardized tools such as the System Usability Scale (SUS) and NASA-TLX. Second, integrating procedural simulation modules—such as hysteroscopic navigation and biopsy training—will expand the application from visualization to active skill acquisition. Third, incorporating multi-modality data (e.g., MRI or CT) alongside US could enrich spatial context and improve cross-modality comprehension. Finally, advances in lightweight neural rendering and real-time segmentation could enable adaptive model updates directly on standalone devices, streamlining data preparation workflows.

5.6. Implications

This work underscores the feasibility of deploying advanced medical visualization systems on accessible standalone VR hardware. By uniting volumetric, surface-based, and population-level data representations, the system offers a scalable framework for both medical education and research into anatomical variability. More broadly, it demonstrates how open-source, reproducible pipelines—spanning 3D Slicer, Blender, and Unity—can be leveraged to translate complex medical data into pedagogically meaningful immersive experiences, ultimately supporting broader adoption of XR tools in healthcare edu

6. Conclusion

This work introduced a standalone VR system for interactive visualization of uterine US volumes, integrating surface meshes, deviation heatmaps, and orthogonal slice exploration within a unified environment. By leveraging an open-source pipeline that combines segmentation, mesh processing, and immersive rendering, we demonstrated that clinically relevant visualizations can be delivered in real time on accessible hardware such as Meta Quest 2. The system enables learners and clinicians to bridge conventional 2D imaging with immersive 3D exploration, fostering improved spatial understanding of uterine morphology and anatomical variability.

Preliminary demonstrations highlight the educational potential of combining volumetric intensity data with population-based shape analysis, particularly for training scenarios where exposure to real patients is limited. At the same time, limitations remain in terms of dataset size, segmentation accuracy, and the absence of structured clinical evaluation. Future work will focus on extending the framework with procedure-specific simulation modules, larger-scale user studies, and integration of multimodal imaging data.

Overall, the results underline the feasibility and promise of deploying immersive medical visualization on lightweight, standalone hardware. Such systems have the potential to democratize access to advanced training tools, reduce reliance on tethered setups, and open new avenues for research and education in gynecology and beyond.

References

[1] F. Pires, C. Costa, P. Dias, On the use of virtual reality for medical imaging visualization, Journal of Imaging 4 (2021) 1034–1048. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC8455774. doi:10.1007/s10278-021-00480-z.

[2] J. Yuan, S. S. Hassan, J. Wu, C. R. Koger, R. R. S. Packard, F. Shi, B. Fei, Y. Ding, Extended reality for biomedicine, Nature Reviews Methods Primers 3 (2023) 2662–8449. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC10088349. doi:10.1038/s43586-023-00208-z.

[3] Q. Liu, S. Qiu, Y. Wang, et al., Coordinated 2d-3d visualization of volumetric medical data in xr with multimodal interactions, 2025. URL: https://arxiv.org/html/2506.22926v1. arXiv:2506.22926.

[4] M. Lang, S. Ghandour, B. Rikard, E. K. Balasalle, M. R. Rouhezamin, H. Zhang, R. N. Uppot, Medical extended reality for radiology education and training, Journal of the American College of Radiology 21 (2024) 1583–1594. doi:https://doi.org/10.1016/j.jacr.2024.05.006.

[5] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, 2015, pp. 234–241. URL: https://doi.org/10.1007/978-3-319-24574-4_28. doi:10.1007/978-3-319-24574-4_28.

[6] F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, K. H. Maier-Hein, nnu-net: a self-configuring method for deep learning-based biomedical image segmentation, Nat Methods (2021). URL: https://doi.org/10.1038/s41592-020-01008-z. doi:10.1038/s41592-020-01008-z.

[7] Q. Jin, Z. Meng, C. Sun, H. Cui, R. Su, Ra-unet: A hybrid deep attention-aware network to extract liver and tumor in ct scans, Frontiers in Bioengineering and Biotechnology Volume 8 - 2020 (2020) 156–165. URL: https://doi.org/10.3389/fbioe.2020.605132. doi:10.3389/fbioe.2020.605132.

[8] X. Yang, H. Dou, R. Li, X. Wang, C. Bian, S. Li, D. Ni, P.-A. Heng, Generalizing deep models for ultrasound image segmentation, in: Medical Image Computing and Computer Assisted Intervention – MICCAI 2018, Springer International Publishing, 2018, pp. 497–505. doi:10.1007/978-3-030-00937-3_57.

[9] N. Tajbakhsh, L. Jeyaseelan, Q. Li, J. N. Chiang, Z. Wu, X. Ding, Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation, Medical Image Analysis 63 (2020) 101693. doi:10.1016/j.media.2020.101693.

[10] E. Boneš, C. Bohak, M. Gergolet, Žiga Lesar, M. Marolt, Automatic segmentation and alignment of uterine shapes from 3d ultrasound data, Computers in Biology and Medicine 178 (2024) 108794. URL: https://doi.org/10.1016/j.compbiomed.2024.108794. doi:10.1016/j.compbiomed.2024.108794.

[11] Žiga Kokelj, C. Bohak, M. Marolt, A web-based virtual reality environment for medical visualization, in: 2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO), IEEE, 2018, pp. 0299–0302. URL: https://ieeexplore.ieee.org/document/8400057. doi:https://doi.org/10.23919/MIPRO.2018.8400057.

[12] Žiga Kokelj, Implementacija navidezne resničnosti v spletno vizualizacijsko ogrodje Med3D, Master's thesis, Fakulteta za računalništvo in informatiko, Univerza v Ljubljani, 2018. URL: https://core.ac.uk/download/pdf/162020494.pdf.

[13] F. Moro, M. T. Giudice, M. Ciancia, D. Zace, G. Baldassari, M. Vagni, H. E. Tran, G. Scambia, A. C. Testa, Application of artificial intelligence to ultrasound imaging for benign gynecological disorders: systematic review, Ultrasound in Obstetrics & Gynecology (2025) 295–302. URL: https://doi.org/10.1002/uog.29171. doi:10.1002/uog.29171.