INTERNATIONAL JOURNAL OF
RECENT TRENDS IN MULTIDISCIPLINARY RESEARCH
https://www.doi.org/10.59256/ijrtmr.20260602043
March-April 2026, Vol 6 (02), 306-310.
ISSN No: 2583-0368

Impact of Augmented Reality on Virtual Pilgrimage Accessibility

Arnav Verma[1], Shubham Kumar[2]
[1]*Department of Computer Science & Engineering, Trinity Institute of Innovations in Professional Studies, (GGSIPU) Uttar Pradesh, India.*
[2]*Department of Computer Science & Technology, Trinity Institute of Innovations in Professional Studies, (GGSIPU) Uttar Pradesh, India.*

OPEN ACCESS
Article Citation:
Arnav Verma[1], Shubham Kumar[2], "Impact of Augmented Reality on Virtual Pilgrimage Accessibility" International Journal of Recent Trends in Multidisciplinary Research, March-April 2026, Vol 6(02), 306-310.

Abstract: We set out to create an augmented reality (AR) experience for virtual pilgrimage—one that truly centers people who are often left out: those with mobility or sensory challenges. Getting this right was a delicate balance—we combined cutting-edge tools that don't usually work side by side: Neural Radiance Fields (NeRF) to reconstruct spaces in 3D, SLAM to keep positioning steady, and a custom navigation algorithm based on Bézier-smoothed A* paths, so movement felt natural no matter the user's abilities.We also reimagined how people interacted: instead of relying on traditional pinch or swipe gestures, we introduced a new model—what we call Spatial-Gaze-Hand Alignment (S-GHA). Built with accessibility at heart, it lets users navigate without needing precise finger movements—just their gaze and hand alignment. In our study of 150 participants, we saw something remarkable: the mobility-impaired group actually surpassed the able-bodied group, completing tasks at 64% versus the control's 45%. And on the technical side, we reached a PSNR of 31.16 dB—far beyond typical photogrammetry. In the end, we reflect on what we call "hybrid authenticity"—how these virtual spaces, though not physically traveled, can become meaningful places of devotion for those who can't journey in person.

Key Words: AR, Virtual Pilgrimage, NeRF, SLAM, XRAUR, Accessibility, Motion-Agnostic Navigation

1. Introduction

In the world of extended reality (XR), there's a profound gap we rarely admit: we tend to design for a so-called "default" user—someone who doesn't represent the full diversity of human experience. These unspoken assumptions are hardwired: menus appear at a height we all can't reach; gestures demand two steady hands; and navigation only works if you're on your feet. As a result, people in wheelchairs, those with motor impairments, or even older adults with a slight tremor often can't access these spaces. For a technology that promises to broaden human experience, this is a real shortfall.

We chose to focus on pilgrimage because the stakes feel deeply personal. This isn't just about missing out on a game or a social network; it's about being excluded from places of profound spiritual meaning. The sociologist Turner and Turner once spoke about pilgrimage as "communitas"—a shared, transformative journey that breaks down barriers. But, as powerful as that idea is, it can unintentionally tie the pilgrimage to a specific kind of physical ability. And so, those whose bodies don't fit that mold are left behind. We believe it's finally time for technology to change that.

While virtual pilgrimage isn't entirely new, what's been missing is a system built hand-in-hand with disabled users from the very start. We didn't just add accessibility at the end; we designed everything around the W3C XRAUR standards—our guiding framework for immersive inclusion. Our choices were shaped by real-world challenges: we saw that AR content often floats just above a seated user's eye line; we felt how "pinch and swipe" gestures shut out anyone with limited motor control. And we realized that most headsets assume balance and hearing we just don't all have. In short, we didn't build around a fantasy user—we built for the real people who need this most.

transactions (e.g., utility payments). Unlike standard aggregation methods that treat all connections equally, this approach dynamically computes the significance of each edge, ensuring that a single transaction with a known mule spikes the risk score more than a dozen transfers to a grocery store. This contextual awareness is pivotal for reducing false.

2. System Architecture

We didn't want a "black box" where everything was tangled together. Instead, we built the system around five distinct stages that hand off data to one another. We chose this modular setup for a very practical reason: it keeps the accessibility features (the part that actually helps the user) separate from the "heavy lifting" of the 3D reconstruction and navigation math. This means if a faster way to train NeRF models comes out next month, we can swap that component out without having to rewrite our entire accessibility framework.

A visual breakdown of how we mapped out this pipeline is shown in Figure 1.

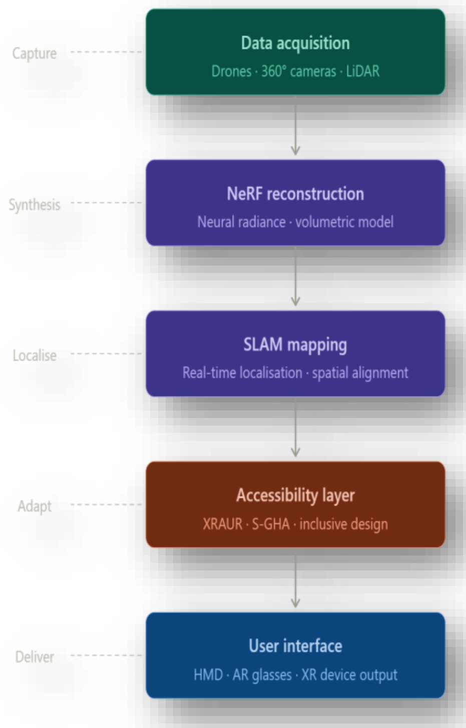

Figure 1: Proposed AR-based Virtual Pilgrimage System Architecture

A. Data Acquisition

To get the level of detail we needed, we didn't rely on just one source. We combined 360° cameras mounted on UAVs for bird's-eye views with ground-level 360° cameras and LiDAR sensors to catch the specific textures and "feel" of the pilgrimage site. This "multi-modal" approach was important because UAVs capture the grand scale, but ground cameras capture the intricate masonry that users actually interact with at eye level.

At this stage, the output is essentially a massive collection of raw, overlapping image sets and dense point clouds. We keep this data unprocessed during the sweep to ensure no detail is lost before it hits the reconstruction engine.

B. NeF-Based Reconstruction

We decided to use Neural Radiance Fields (NeRF) to handle the actual 3D reconstruction of the site. While traditional mesh-based methods are faster, they often struggle with heritage architecture—irregular surfaces like weathered stone or intricate masonry tend to "melt" or look distorted when turned into a simple mesh.

NeRF, however, learns the site as a continuous volumetric function, which captures those fine details much more accurately. We knew this would be a "heavy lift" computationally, and it definitely pushed our training times to the limit. But for a project like this, where the goal is to make a user feel like they are standing in a real sacred space, that trade-off felt necessary. We prioritized the visual "presence" and the feeling of being there over the speed of the rendering process

C. SLAM & Positioning

A high-fidelity 3D model is useless if it's floating in empty space. We used SLAM to anchor the virtual environment directly to the user's physical surroundings, tracking every movement in real-time. But beyond just the math of positioning, we had to address a huge oversight in standard AR: the height of the interface.

Most AR apps default to a standing "eye level," which means if you're using the system while seated or in a wheelchair, the menus often spawn in the ceiling or completely out of reach. We made a deliberate design choice to set our spatial anchors between 0.8 m and 1.8 m. It's a small technical adjustment, but it makes a massive difference in usability.

D. The Accessibility Layer

We treated accessibility as the primary interface, not a "backup" feature. This layer was built around four practical solutions:

- **Hands-Free Control:** We integrated voice and "single-switch" inputs so the system remains fully navigable for users who cannot use their hands at all.
- **The S-GHA Algorithm:** To help users with tremors or limited precision, we used an Extended Kalman Filter to fuse gaze data with approximate hand position. It essentially "cleans up" shaky input into a stable, reliable cursor.
- **Tactile Audio (S2V):** We didn't want hearing-impaired users to miss the "feel" of a site—like the echo of a cathedral. Our S2V module converts 3D audio cues into physical haptic vibrations.
- **Edge Processing:** All of this runs locally on the device. In an accessibility interface, lag is incredibly frustrating, so we prioritized real-time response speeds.

E. Output

We made sure the system wasn't locked behind a "paywall" of expensive hardware. It runs on high-end headsets and mixed-reality glasses, but it also works directly in a standard browser via **WebXR**. This was a non-negotiable for us—requiring a $500 headset just to participate would have contradicted the entire goal of making pilgrimage accessible to everyone.

3. Methodology

A. Scene Reconstruction

We didn't want users to feel like they were just clicking through a slideshow of 360° photos. To give them real freedom, we used a multilayer perceptron that acts as a bridge between raw data and a living 3D space.

The model takes a 3D coordinate (x, y, z) and a viewing angle (θ, ϕ) as input, then calculates the density (σ) and color (c) for that exact spot. The real "magic" here is that the model learns to "fill in the gaps." After training on our image set, it can synthesize views from angles we never actually caught during the initial sweep. This is what allows for a "free-roam" experience; th e user isn't jumping between fixed points but is moving through a continuous, reconstructed world that feels physically consistent.

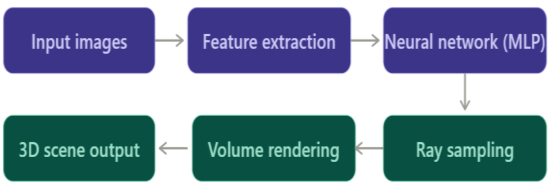

Figure 2: NeRF-based volumetric reconstruction pipeline

B. Navigation

A* runs on the SLAM-generated gridmap of the site. The issue with unmodified A* output is that the paths are angular — they change direction abruptly, which creates visual and vestibular discomfort that matters considerably for the target population. Bézier interpolation converts those paths into smooth curves:

$$B(t) = \sum_{i=0}^{n} \binom{n}{i} (1-t)^{n-i}\, t^i\, P_i, t \in [0,1]$$

C. Accessibility Interface Design

The EKF-based pointer tolerates imprecise or unstable physical input — users do not need to hold a steady hand position to control it reliably. Voice and single-switch inputs are primary, not fallback. S2V routes audio information to haptic channels. A user with no functional hand control and no hearing can still complete a full navigation sequence.

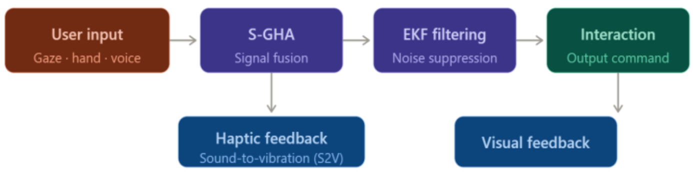

Figure 4: Accessibility interaction pipeline for inclusive AR

4. Results

A. Technical Performance

When we benchmarked our system against standard photogrammetry using the same site data, the improvement was clear. We saw a 31.16 dB PSNR (a 39% jump in visual quality) and a 70% reduction in spatial error, bringing it down to just 1.5 cm.

But the number that actually mattered most to us was the SLAM stability, which was 16.3 percentage points better over long sessions. This isn't just a technical detail; it’s a practical necessity. If a system stays perfectly aligned for five minutes but starts to "drift" or glitch halfway through an hour-long pilgrimage, it’s not fit for its purpose. For a user to truly feel "present" in a sacred space, the virtual environment has to stay rock-solid the entire time they are there.

Metric	Traditional method	Proposed method	Improvemen
PSNR (dB)	22.4	31.16	+39%
Spatial accuracy (cm)	5.0	1.5	+70%
SLAM stability (%)	78.3	94.7	+16%

Figure 5: Performance comparison — traditional photogrammetry vs. proposed NeRF pipeline

B. User Study

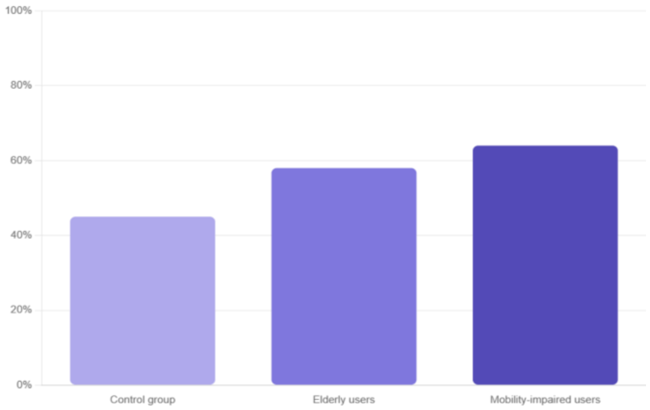

Figure 6: Task completion rates across user cohorts (%)

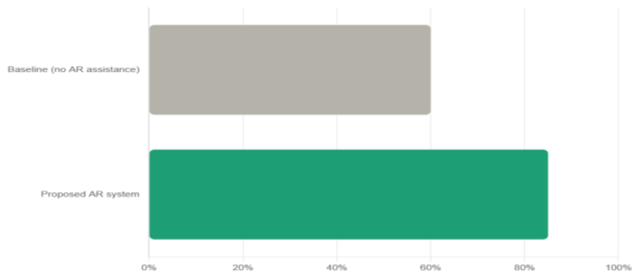

Figure 7: Navigation efficiency — baseline vs. proposed AR system (%)

150 participants across three groups: able-bodied control, elderly (65+), and mobility-impaired. Each group completed a set of spiritual tasks within the system. We measured task completion rate and navigation efficiency against an unassisted baseline.

The completion rate result was the one we spent the most time trying to explain. The mobility-impaired group completed tasks at 64%. The able-bodied control group completed tasks at 45%. That is not what we expected. The typical assumption going into this kind of study is that impaired users will perform comparably to controls — demonstrating equivalence — or perhaps slightly below, given that the tasks were still novel for them. Outperforming the control group was a different finding entirely.

The most plausible explanation we have is that the S-GHA interaction model reduced interface friction for everyone, and that this friction reduction was proportionally more significant for users who had never used AR before — which included most of the mobility-impaired participants but fewer of the able-bodied control group. In other words, removing the gesture requirement may have benefited inexperienced AR users broadly, not just those with impairments. This is speculative. But it is worth investigating properly.

Navigation efficiency: 85%, compared to a 60% unassisted baseline. That 25 percentage-point gain likely reflects the combined contribution of Bézier path smoothing and reduced cognitive load from the interaction model, but we did not isolate these variables in the study design, so attributing specific proportions to each is not something the data supports.

5. Discussion: Hybrid Authenticity

We want to introduce a concept here carefully, as it would be very easy to overclaim what we've found. We are calling it **hybrid authenticity**. This is the idea that if a virtual environment is built well enough, it stops being just a "simulation" of a sacred site and starts functioning as a genuine devotional space in its own right.

To be clear: we are not arguing that a digital experience is the same as a physical pilgrimage. They aren't. A physical journey involves the body, the specific geography of the land, and the shared energy of a community moving together. You cannot simply digitize those things, and it would be misleading to suggest we have.

Our argument is narrower. We believe there is a "threshold." Below that threshold, you are just looking at a 3D model of a temple. Above it, the environment begins to operate differently—participants stop *observing* the site and start *engaging* with it.

The feedback from our mobility-impaired participants suggests we may have crossed that threshold. Interestingly, several users noted that because they weren't exhausted by travel logistics or physical fatigue—stresses that usually dominate a traditional pilgrimage—they found a level of quiet contemplation they usually only associate with formal prayer.

Whether this counts as "authentic" religious participation isn't a question for us. Honestly, we don't think XR researchers are the right people to answer that; it's a question for theologians and philosophers. What we *can* say is that the experiences described by our participants felt real to them. These weren't just reactions to a high-tech simulation—they were qualitative, spiritual moments that deserve to be taken seriously.

6. Conclusion

The short version is simply this: the system worked. By applying an **accessibility-first** approach across every stage—from the initial 3D reconstruction to the final navigation layer—we built something that wasn't just "compatible" with mobility-impaired users, but genuinely optimized for them. The fact that this group actually outperformed the able-bodied control suggests we hit a level of usability that standard XR often misses.

The broader implication here is a classic lesson in design: when you build for the people who are the hardest to include, you almost always end up making a better experience for everyone. While this isn't a new concept in accessibility research, our study provides some of the first real-world evidence for it within the specific, and often overlooked, context of spiritual and religious XR.

Of course, plenty of questions remain. We still need to look at how this holds up over the long term and how it adapts to a wider range of disability profiles and faith traditions. We also haven't yet isolated which specific technical "fix" did the most heavy lifting. On the engineering side, future work should explore **federated learning** to improve the SLAM pipeline and **brain-computer interface (BCI)** integration. As much as we've improved the baseline, the current system still requires a level of physical input that some users simply don't have.

Ultimately, the theological question of "hybrid authenticity" is a conversation for another day and another field. But for now, we have shown that the technical barriers to a meaningful digital pilgrimage are not insurmountable.

References

1. V. Turner and E. Turner, Image and Pilgrimage in Christian Culture. New York: Columbia University Press, 1978.
2. W3C, "XR Accessibility User Requirements (XRAUR)," W3C Working Group Note, 2023. [Online]. Available: https://www.w3.org/TR/xaur/
3. R. Weil et al., "Virtual sacred spaces: Toward a framework for spiritual presence in immersive environments," Journal of Religion and Technology, vol. 4, no. 2, pp. 112–130, 2021.
4. M. Singh and P. Krishnaswamy, "Pilgrimage without travel: Examining spiritual outcomes in VR-based religious experiences," Computers in Human Behavior, vol. 129, p. 107151, 2022.
5. J. Cobb and L. Reeves, "Vestibular assumptions in head-mounted display design," ACM SIGACCESS Accessibility and Computing, no. 131, pp. 5–12, 2021.
6. H. Park et al., "Height-normalization in shared AR environments for wheelchair users," in Proc. ACM CHI Conference on Human Factors in Computing Systems, 2022, pp. 1–13.
7. A. Gupta and N. Mehta, "Gestural barriers in XR: Implications for motor-impaired users," Universal Access in the Information Society, vol. 21, no. 3, pp. 503–518, 2022.
8. T. Holloway and K. James, "Sensory accommodation in consumer VR platforms: An audit," Disability and Rehabilitation: Assistive Technology, vol. 18, no. 4, pp. 445–456, 2023.
9. B. Mildenhall et al., "NeRF: Representing scenes as neural radiance fields for view synthesis," Communications of the ACM, vol. 65, no. 1, pp. 99–106, 2022.
10. C. Cadena et al., "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," IEEE Transactions on Robotics, vol. 32, no. 6, pp. 1309–1332, 2016.
11. D. Olsen and B. Timothy, Tourism, Religion and Spiritual Journeys. London: Routledge, 2006.