https://doi.org/10.2352/EI.2026.38.13.ERVR-188

Virtual Reality-Based AI Mental Health Companion: A Multimodal System for Therapy, Mindfulness, and Stress Detection

Sharad Sharma, Johannas Katikala, Keerthana Srinivasan, Pretom Roy Ovi, Department of Data Science, University of North Texas, Denton, Texas 76207, USA
Sam Rahimzadeh Holagh, Navatha Maloth, Ana Cleveland, Department of Information Science, University of North Texas, Denton, Texas 76207, USA
Timothy McMahan, Regina Kaplan-Rakowski, Department of Learning Technologies, University of North Texas, Denton, Texas 76207, USA

Abstract
Virtual Reality (VR) has recently attracted more attention in mental health applications due to its ability to immerse users in controlled and interactive environments. This paper presents a VR-Based AI Mental Health Companion, a multimodal system designed to support therapy, mindfulness, and real-time stress detection within immersive virtual reality (VR) environments. The system integrates artificial intelligence (AI) techniques by including natural language processing, emotion recognition, and physiological signal analysis by creating personalized mindfulness experiences and interactive meditation coaching. GPT-powered non-player characters (NPCs) are designed with specific therapeutic roles in mind, such as guided mindfulness facilitation, emotional support, and stress-aware conversational therapy. The system employs pose estimation to identify key body points and apply rule-based logic to assess posture accuracy during guided yoga exercises, providing real-time feedback to support correct movement execution. The work also includes biometric integration, such as EEG monitoring for enhanced emotional sensing. Expanding language support, increasing the diversity of pose datasets, and incorporating feedback from clinical professionals help refine the system. By combining immersive VR environments, GPT-driven therapeutic NPCs, and real-time posture validation, the proposed VR-Based AI Mental Health Companion demonstrates the potential of AI–VR convergence as a scalable approach to mental health care, with promising applications in stress management and preventative therapy.

1. Introduction
The increasing prevalence of mental health concerns worldwide has prompted researchers and developers to explore innovative ways to deliver effective, scalable, and personalized mental health interventions. One promising direction is the use of immersive technologies such as Virtual Reality (VR), which offers a highly engaging and controlled environment for therapy, mindfulness, and emotional support. Virtual Reality Therapy (VRT) has been clinically validated for addressing a range of psychological disorders, including anxiety, post-traumatic stress disorder (PTSD), and phobias [1,2]. Unlike traditional digital therapy tools, VR systems offer an immersive context where individuals can interact with lifelike environments and agents, allowing for heightened presence and emotional engagement.

Recent advancements in generative artificial intelligence (AI), especially large language models (LLMs), have added a new dimension to the capabilities of NPCs (non-player characters) within these environments. These AI agents, powered by models like GPT, enable natural language understanding and generation, which allows users to engage in real-time, emotionally intelligent conversations. Such interactions can provide not only informational content, but also empathetic responses, personalized interventions, and dynamic adaptations based on the user's mood and context [3,4]. The convergence of VR and generative AI opens the door to creating emotionally aware virtual companions that simulate therapeutic interactions.

This study presents a novel Virtual Reality Instructional (VRI) environment developed in Unity, where LLM-powered NPCs offer supportive conversations and guide users through various mental health exercises. The system includes multiple modalities such as a hospital-themed 3D environment, NPCs for cognitive-behavioral therapy (CBT) and mindfulness, an anxiety assessment module, and a pose estimation component using YOLOv8 and OpenCV. This multimodal setup not only enables emotional dialogue but also physical feedback via posture analysis during stress-relief activities such as yoga. Through this integration, the system aims to offer a holistic digital companion for emotional and psychological support.

Figure 1. VR-Based AI Mental Health Assistant Using ConvAI

Figure 1 shows our developed AI Assistant using ConvAI and VR by incorporating Meta Quest 3 Touch controllers for course of action, visualization, and situational awareness. The rest of the paper is structured as follows. Section 2 briefly describes the work done previously. Section 3 describes the designed VR environment; Section 4 describes the implementation of the VR-based mental health companion and the pose estimation implementation in VR, Section 5 discusses the Multimodal System for Therapy, EEG monitoring and Stress Detection, and Section 6 lists the conclusion and future work.

2. Related Work

Human pose estimation has gained significant attention in computer vision due to its applications in healthcare, human-computer interaction, surveillance, and virtual environments. In recent years, the YOLO family of models has emerged as a popular object detection framework, offering real-time performance while maintaining accuracy [5],[6]. The integration of YOLO into pose estimation pipelines has been widely explored, particularly in fall detection, gesture recognition, and immersive simulations [7], [8].

YOLO-based models, such as YOLOv4 combined with key point refinement, provide fine-grained hand pose estimation, supporting effective gesture control in complex environments [9]. YOLOv8's real-time key point detection facilitates applications like yoga pose correction and physical therapy monitoring by providing accurate posture feedback [3]. Human posture tracking with YOLO improves video surveillance systems by enabling the detection of suspicious behaviors in real time. Integration of YOLO with depth sensors such as Kinect improves fall detection accuracy and robustness in elderly care settings [10]. In a practical application, current study developed a VR-based academic stress and anxiety management system utilizing YOLOv8 pose detection integrated with Unity 3D. The system tracked user postures in real-time during relaxation activities and breathing exercises. YOLOv8 extracted key points corresponding to upper body movements, which were analyzed to assess user engagement in VR therapy sessions. This integration enabled dynamic feedback and avatar interaction based on detected posture, enhancing immersion and supporting emotional regulation. Experimental results demonstrated promising accuracy and low latency, indicating the approach's suitability for responsive biofeedback environments.

3. VR Environment for Meditative Experiences.

At the heart of the experience is a 3D hospital model constructed in Unity (refer Figure 2). This environment features various rooms and stations designed for different mental health applications, including therapy, meditation, and daily check-ins. The spatial layout and environmental design play a crucial role in setting the tone for the interaction, creating a calming, non-threatening atmosphere conducive to emotional openness. To enhance visual fidelity and optimize performance across a variety of VR-capable hardware, the system employs the Universal Render Pipeline (URP). URP enables lightweight rendering while still supporting advanced effects such as real-time shadows and post-processing.

Figure 2. VR Wellness Center Environment (or UNT Clinic)

In terms of lighting design, baked lighting was used extensively throughout the hospital model to simulate realistic illumination with minimal computational overhead during runtime. Baked lighting allows for precomputed lightmaps, enhancing both the realism and performance of the environment. Real-time lights were selectively used for interactive elements, such as moving NPCs or dynamic visual cues. Reflection probes and light probes were also implemented to ensure smooth and immersive transitions between different lighting zones and to support dynamic objects moving through the environment.

Figure 3 VR Meditation Environment

As seen in Figure 3, the meditation environment is designed as a nature-immersive virtual space inspired by a jungle setting with waterfalls, dense greenery, and softly filtered natural light. The designed environment is intentionally aimed at shifting the focus of clinical aesthetics towards the principles of biophilic design which promote relaxation and emotional control among the students. There is ambient nature sound, such as running water and far birds, which go on throughout the session to induce relaxation and help the participant become attentional to factors other than academic pressures. Collectively, the visual immersion and sound that is spatialized improves the elements of presence and psychological safety, which are significant requirements to alleviate the anxiety-inducing factor and promote restorative mental conditions in learning settings. The yoga exercise room (refer Figure 4) is a structured, yet supportive VR environment dedicated to physical mindfulness and movement-based stress relief. In this space, animated virtual characters demonstrate yoga poses and movement sequences in real time, serving as visual references for learners. Users are encouraged to follow these movements while the system continuously monitors their posture using image processing techniques based on YOLO pose estimation.

Figure 4 VR Yoga Exercice Room Environnent

The system extracts key body joints and compares the user's movements against predefined pose rules to assess accuracy and alignment. Based on this analysis, real-time feedback is provided either correcting the user's posture, offering guidance for improvement, or confirming correct execution before advancing to the next pose. This closed-loop interaction transforms yoga practice into an adaptive learning experience, allowing students to safely learn, self-correct, and progress at their own pace while maintaining engagement and physical awareness within the VR environment. Embedded within this environment are ConvAi-powered GPT NPCs (refer Figure 5) that serve as the primary interface for user interaction. These agents engage in real-time conversations with users, helping them reflect on their emotional state, explore cognitive-behavioral techniques, and participate in mindfulness activities. The NPCs are capable of mood detection, enabling them to adapt their responses based on the user's current affective state. This makes the interactions more authentic and supportive, mimicking the empathy and attentiveness of human therapists [11, 12].

Figure 5 ConvAi NPC Therapist

The camera and IMU components work together to update the user's location indicator in the Mini Map which reflects the 3D virtual environment. The user then interacts with the interface via buttons and dropdown menus. The user selects the appropriate floor, wing letter, and room number of their desired location. Alternatively, the user can choose to find the nearest building exit. Based on their input, the system designates a target as the end destination and displays appropriate route information to the user. The MARA generates the navigation path in both the virtual and physical environment.

4. VR-Based AI Mental Health Companion

The VR Based AI Mental Health Companion will be applied in the Virtual Reality Instructional environment as a facilitated workflow integrating conversational assistance with immersive interventions. GPT-driven NPCs are implemented based on the ConvAI system and given specific therapeutic stresses, such as virtual therapist style guidance, meditation facilitation, coaching stress relief, and companionship. Prompt engineering is employed to make sure that the response is context relevant and encouraging over a broad input of users. As an illustration, in the case the user is anxious or distressed, the NPC should react by providing reassuring messages, cognitive reframing propositions, and willingness to activate the short-term mindfulness process (grounding exercises, breathing exercises, etc.).

One design objective of the companion is emotional intelligence. Sentiment cues on speech are used by the NPC to determine affective state and modify responses based on them, based on current directions of interactive agents being dynamically responsive to the psychological needs of the user, like in MIND [5]. To ensure additional consistency as well as minimize the conversational drift between interactions, the assistant is trained with a fixed behavioral profile governing the response style. This assistant set up is described in fig. 6 using a trait-based profile of Openness, Meticulousness, Extraversion, Agreeableness, and Sensitivity. It is possible to fix these parameters, to maintain a predictable and supportive tone and to enable the NPC to adjust the contents of its replies as per the results of the questionnaire seizures and the context of the current dialogue with the user. Moreover, the session level memory systems enable personalization, since NPC can remember user preferences and previously chosen intervention which enhances continuity and trust in the long run [3]. User engagement is strengthened by the system's ability to feel emotionally responsive rather than generic, which is an important factor in digital mental health effectiveness [13].

4.1. Anxiety Assessment

The other key part of the companion is the anxiety assessment module, which is presented as a structured questionnaire in the form of a conversation with the NPC in the VR clinic. The agent uses dialogue format to ask standardized questions, which he records answers to, and interprets the outcomes to prescribe an effective therapeutic activity. The NPC is also able to alter follow up prompts on the reactions of the users such as clarifying questions in case of answers that show high stress or uncertainty. In case the assessment shows that there is more anxiety, the system can suggest entering the meditation room and being guided into relaxation, and those users who require more active coping suggestions can be directed to the yoga room to be able to use the movement-based stress relievers. Administering the questionnaire in a conversational format decreases the clinical aloofness of standard survey interfaces and may enhance the user experience and accuracy of the responses, which is consistent with the evidence that user customized conversational interventions can enhance engagement and outcomes in online mental health environments [7].

4.2. Pose Detection and Feedback Mechanism

In addition to the support that is provided by dialogue, the system uses a physical feedback system by using real time pose estimation to enhance embodied mindfulness in the process of performing yoga-based relaxation. The system identifies 17 body key points and produces a skeletal reconstruction of the user with the help of YOLOv8 and OpenCV. Rule based posture validation is applied to these key points and comparisons are made between geometric relationships between joint angles and limb alignments and pose specific constraints. This strategy was chosen due to the need to be responsive in VR in real time and eliminate the data and training needs of pose classification models that need large yoga labeled datasets.

Figure 7 shows qualitative illustrations of the adopted pose pipeline. The images display the identified user area with a bounding box and confidence rate, and the skeleton key point generated in real time. The examples show that the key point tracking under varying indoor environments is stable, which proves the feasibility of the application of the same set of rules in different backgrounds and lighting conditions. The rule-based logic is used during yoga sessions whereby it makes use of these key points to give immediate feedback whether the posture is performed correctly or the user needs to reposition themselves before continuing the session.

The proposed system has been tested with various types of users with different body types and fitness levels. This has helped in increasing the overall generalization of the rule-based constraints. A two-stage approach using YOLO for keypoint extraction and an SVM classifier has also been attempted; however, this has not yielded significant accuracy for real-time corrective feedback with varying pose executions. Moreover, retraining the pose model for yoga classes has also been attempted; however, this has not been feasible within the given time frame and resources [6,14]. Instead, the proposed system focuses on interpretability and low-latency operation with potential for extension with more pose data in the future.

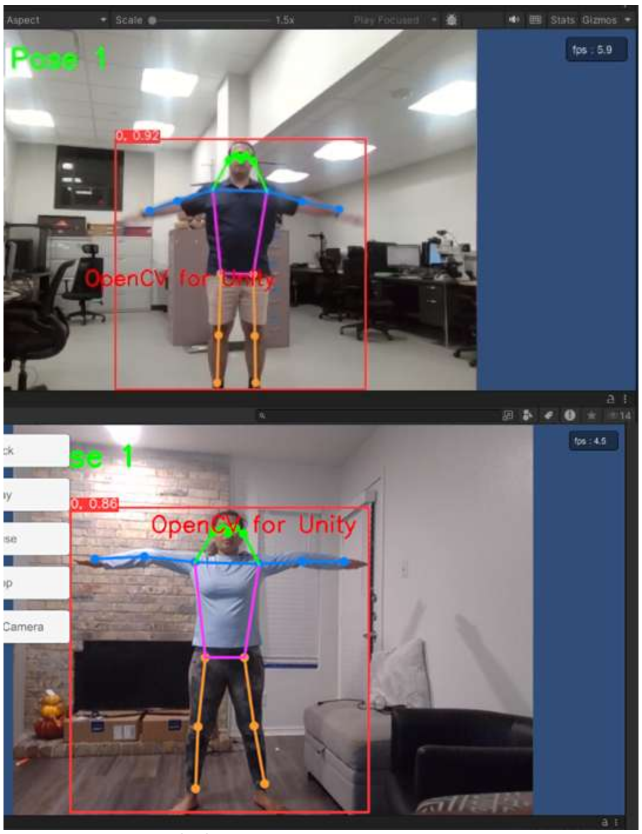

Figure 7 Real time YOLOv8 keypoint based pose detection in Unity using OpenCV, showing bounding box confidence and 17 keypoint skeleton overlays across two indoor environments

5. Multimodal System for Therapy, EEG Monitoring, and Stress Detection

This section has highlighted the EEG monitoring layer of the VR-based mental health companion for real-time tracking of the cognitive state during the immersive experience. This has been designed for the real-time observation of the fluctuations in the mental state of the participant while interacting within the virtual world. Emotiv EEG has been incorporated for the real-time visualization, verification, and logging of stress-related metrics. Figure 8 depicts live attention, engagement, stress, relaxation, and excitement values streamed from the Cortex MET output during the VR session.

5.1. Emotiv EEG Streaming and Unity Integration

The acquisition of EEG data was achieved through the Emotiv Cortex WebSocket API. The Unity client requested access, authorized the application, retrieved connected headsets, opened a session, and subscribed to live data streams. WebSocket messages were processed asynchronously before being dispatched to Unity's main thread to maintain UI stability during runtime in VR.

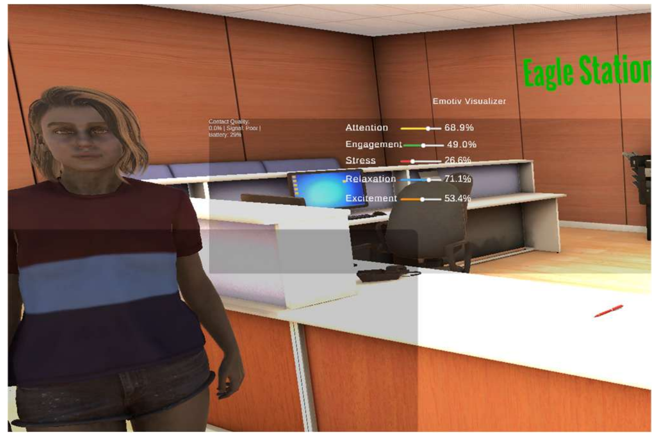

Figure 8 Emotiv metrics overlay in Unity.

5.2. Streams Used for Monitoring and Reliability Checking

Three Cortex streams are subscribed for monitoring and quality control. The Mental Performance Metrics stream was used for EEG-derived cognitive state indicators for attention, engagement, stress, relaxation, and excitement. The device stream was used for monitoring headset status, i.e., battery percentage and contact quality. The system stream was used for monitoring integrity of EEG signals by using signal quality events that provide quality values for each channel. These auxiliary diagnostics helped in understanding changes in cognitive metrics and reduced the chances of misinterpreting decreased contact as stress.

5.3. Two Implemented Methods for Stress and Cognitive State Estimation

There were two techniques used in estimating stress and indicators of cognitive states during VR sessions.

The former applied Cortex derived mental performance scores. In this method, stress and attention were not calculated on the basis of raw EEG waves that Unity made. Rather, it devoured the normalized values of state of the mind generated by Cortex and mapped them to named variables which were used in the VR workflow. The met payload was decoded into a numeric array in order and transformed into a dictionary of metrics through a fixed index mapping. Met index 1, engagement index 3, excitement index 5, stress index 7 and relaxation index 9 were all extracted to get attention, engagement, excitement, stress, and relaxation respectively. These values were normalized and held constant between 0 and 1 and represented as a percentage in the monitoring interface by multiplying with 100.

The second technique estimated the proxies of cognitive states in Unity, through EEG spectral features. EEG or band power was streamed and summarized over sliding windows into standard frequency bands, such as delta (1-4 Hz), theta (4-8 Hz), alpha (8-13

Hz) and beta (13-30 Hz). The band power values were also summed up over the chosen channels and transformed to interpretable indices using ratio-based formulations. The proxy of the stress/arousal was estimated by a beta to alpha relationship of the form Stress Raw = betaPower/(alphaPower plus a small constant). Attention and mental workload were calculated through a composite relationship which rose when theta and beta prevailed over the alpha and was calculated as Attention = (thetaPower + betaPower)/(alphaPower + a constant). The computation of relaxation was based on an alpha dominance relationship and was expressed by Relaxation is alphaPower/(betaPower + constant). Engagement has been calculated by taking a composite index of activation, a summation of the beta and theta values over alpha in a steady fashion in line with prolonged task involvement. Since the features obtained through the EEG were different among the participants, the baseline calibration was done with a short rest of about 30-60 seconds. Each index was calculated to obtain the baseline mean and standard deviation and z normalization was done to convert the obtained raw ratios to standardized scores. The normalized values were then mapped to a range of 0 to 1 to be drawn and averaged out the jitter by smoothing over time. This feature-based design offered transparency and customization but needed to be carefully filtered and controlled on artifacts since motion and muscle activity in VR brought noise, particularly in higher frequencies.

5.4. Real Time Visualization and Monitoring Interface

A user-facing EEG monitoring panel has been implemented, which displays the user's attention, engagement, stress, relaxation, and excitement levels. Each of the above variables has been implemented using a UI slider and a corresponding numeric label. The slider corresponds to the normalized value, and the numeric label corresponds to the value as a percentage. The monitoring panel has been implemented in a manner that is optional, using a toggle interaction.

5.5. Signal Quality Monitoring

The signal integrity was monitored to ensure that cognitive metrics were interpreted under valid conditions. From the dev stream, battery percentage was obtained as well as contact quality, represented as a 0-100 scale with defined thresholds for Excellent, Good, Fair, and Poor. From the sys stream, Signal Quality events were available with per-channel values represented as a discrete 0-4 scale, with lower values indicating better signals. The values were averaged and represented as an inverted percentage measure to generate a total EEG quality estimate. Signal/contact indicators were also represented within the interface to provide a better understanding of cognitive metrics. .

Figure 9. Emotiv headset setup during a Unity session.

Although the feature-based method was integrated into the system, additional processing was needed for preprocessing, baseline calibration, and artifact removal to ensure robustness in real-time VR settings. To ensure accuracy and consistency with the chosen method, stress/cognitive state monitoring was implemented via the Cortex met-based method, where mental performance metrics were calculated internally by Emotiv and streamed live with visualization via Unity

6. Conclusions

The project has demonstrated that a multimodal mental health companion could be created through the combination of virtual reality, generative AI, and computer vision to offer therapy-oriented guidance, mindfulness, and stress reduction. Natural conversation The NPCs generated by the GPT-powered ConvAI can be used to engage in emotionally sensitive and context-aware conversations that can give an intake style questionnaire, recommend, and guide users through the process in a natural manner. The Unity-based clinic setting presupposed a simulated and structured setting in which the users can navigate the special spheres of evaluation, meditation and yoga-oriented relaxation. Also, the YOLOv8 and OpenCV pose estimation module was feasible to give real time confirmation throughout the yoga practice, which is more than conversational support, as the system encourages embodied feedback, which encourages the proper performance of movements. The further work will aim at extending the Emotiv based physiological sensing to more closely support stress sensing, implement language support, diversification and reinforcement of pose rules and datasets, and feedback of clinical professionals to improve safety, usability, and therapeutic relevance. Overall, the results point to the potential possibility of uniting AI and VR and creating immersive adaptive and scalable tools that may be applied as a complement to mental health support under the conditions of education and prevention.

Acknowledgments

This work is funded in part by Sub Award No. NSF00123-08 for NSF Award 2118285.

References

[1] Riva, G., Mantovani, F., Capideville, C. S., Preziosa, A., & Morganti, F. (2007). Affective interactions using virtual reality: The link between presence and emotions. CyberPsychology & Behavior, 10(1), 45–56. https://doi.org/10.1089/cpb.2006.9993

[2] Kandalaft, M. R., Didehbani, N., Krawczyk, D. C., Allen, T. T., & Chapman, S. B. (2013). Virtual reality social cognition training for young adults with high-functioning autism. Journal of Autism and Developmental Disorders, 43(1), 34–44. https://doi.org/10.1007/s10803-012-1544-6

[3] Allison, F., Carter, M., Gibbs, M., & Smith, W. (2018). Design Patterns for Voice Interaction in Games. CHI PLAY 2018 - Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play, 5–17.

[4] Zhang, H., Tsang, K., & Peng, Z. (2025). VChatter: Exploring generative conversational agents for simulating exposure therapy to reduce social anxiety. arXiv

[5] Chen, Y., Li, C., Wang, Y., Xiao, Q., Zhang, N., Kong, Z., & Yan, B. (2025). MIND: Towards immersive psychological healing with multi-agent inner dialogue. arXiv preprint https://arxiv.org/abs/2502.19860

[6] Choudhary, P., Kumar, A., & Raj, A. (2023). Yoga pose detection and feedback generation: A review. International Journal of Multidisciplinary Research and Growth Evaluation

[7] Gaffney, H., Mansell, W., & Tai, S. (2019). Conversational agents in the treatment of mental health problems: Mixed-methods systematic review. JMIR Mental Health, 6(10), e14166. https://doi.org/10.2196/14166

[8] Kandalaft, M. R., Didehbani, N., Krawczyk, D. C., Allen, T. T., & Chapman, S. B. (2013). Virtual reality social cognition training for young adults with high-functioning autism. Journal of Autism and Developmental Disorders, 43(1), 34–44. https://doi.org/10.1007/s10803-012-1544-6

[9] Zhang, H., Tsang, K., & Peng, Z. (2025). VChatter: Exploring generative conversational agents for simulating exposure therapy to reduce social anxiety. arXiv preprint https://arxiv.org/abs/2506.03520

[10] Scapin, S., Echevarría-Guanilo, M. E., Fuculo Junior, P. R. B., Gonçalves, N., & Rocha, P. K. (2018). Virtual reality in the treatment of burn patients: A systematic review. Burns, 44(6), 1405–1416. https://doi.org/10.1016/j.burns.2018.02.011

[11] Cox, S. R., & Ooi, W. T. (2024). Conversational Interactions with NPCs in LLM-Driven Gaming: Guidelines from a Content Analysis of Player Feedback (pp. 167–184), Springer, Cham.

[12] Kenny, P., Parsons, T.D., Gratch, J., Leuski, A., Rizzo, A.A. (2007). Virtual Patients for Clinical Therapist Skills Training. In: Pelachaud, C., Martin, JC., André, E., Chollet, G., Karpouzis, K., Pelé, D. (eds) Intelligent Virtual Agents. IVA 2007. Lecture Notes in Computer Science(), vol 4722. Springer, Berlin, Heidelberg.

[13] Lim, S., Schmälzle, R., & Bente, G. (2024). Artificial social influence via human-embodied AI agent interaction in immersive virtual reality. arXiv

[14] Yeh, S.-C., & Yang, C.-K. (2025). Yoga pose recognition and motion analysis for a home-based fitness monitoring and health management system. Signal, Image and Video Processing.

Author Biographies

Dr. Sharad Sharma is a Sinha endowed Professor in the Department of Data Science and Associate Dean of Research in the College of Information at the University of North Texas (UNT). He is the Director of the Data Visualization and Extreme Reality Lab (DVXR) Lab and the Director of the Data Science & Industry Innovation (DSII) Center at the UNT. He is an expert in modeling and simulation of human behavior for emergency response and decision making with focus on multi-agent systems (MAS), multi-user virtual reality (MUVR), and mobile augmented reality applications (MARA).

Sam Rahimzadeh Holagh is a PhD student in the Department of Information Science at the University of North Texas (UNT). He has worked as a graduate research assistant in the Data Visualization and Extreme Reality Lab (DVXR) Lab at UNT.

Johannas Katikala is a master's student in the Department of Data Science at the University of North Texas (UNT). She has worked as a graduate research assistant in the Data Visualization and Extreme Reality Lab (DVXR) Lab at UNT.

Keerthana Srinivasan is a master's student in the Department of Data Science at the University of North Texas (UNT). She has worked as a graduate research assistant in the Data Visualization and Extreme Reality Lab (DVXR) Lab at UNT.

Navatha Maloth is a master's student in the Department of Information Science at the University of North Texas (UNT). She has worked as a graduate research assistant in the Data Visualization and Extreme Reality Lab (DVXR) Lab at UNT.

Dr. Ana Cleveland is a Regents Professor & Sarah Law Kennerly Endowed Professor in Department of Information Science at UNT. She is also the Director of Health Informatics Program at UNT.

Timothy McMahan is an Assistant Professor in the Department of Learning Technologies and Director of the Mixed Realities Laboratory. His research focuses on Adaptive Virtual Environments. Specifically using neurogaming techniques to build XR environments that adapt to the user to provide a personalized experience. His research also includes the use of video games and simulations for practical applications such as learning and training applications.

Dr. Regina Kaplan-Rakowski is an Assistant Professor in the Department of Learning Technologies at UNT. Her research primarily focuses on Immersive Learning Environments: Virtual Reality, Augmented Reality, Mixed Reality, Virtual Worlds. Her research also includes Online Teaching and Learning, Computer-Assisted Language Learning, and 3D Visualizations of Learning Content.

Dr. Pretom Roy Ovi is an Assistant Professor in Department of Data Science at UNT. His research interest includes enhanced privacy machine learning, federated learning, large language modeling, VR, AR, health informatics