



Deep Learning-Optimized Virtual Reality Interface Design for Anxiety and Depression Treatment: A Patient-Centered Medical Device Innovation

1st Malambo Lushomo *

Zambia Precision Engineering Limited Liability Company

Lusaka, Zambia

malamblo@outlook.com

Received on July 15th, revised on August 14th, accepted on August 24th, published on October 10th

Abstract—Anxiety and depression treatment is constrained by variable adherence and limited personalization in static virtual-reality (VR) protocols. This study developed a patient-centered VR therapeutic system that integrates physiological and behavioral sensing with a deep Q-learning (DQL) adaptive interface. EEG, heart rate, galvanic skin response, gaze, and movement data were converted into a real-time state representation to regulate environmental stimuli within a clinically tolerable exposure range. In a 12-week randomized controlled trial involving 120 participants, Adaptive VRT produced lower week-12 HAMA scores than Non-Adaptive VRT (12.03 ± 5.65 vs. 17.93 ± 5.42 , $p < 0.001$) and lower HAMD scores (12.45 ± 5.85 vs. 15.98 ± 5.33 , $p = 0.002$). The Adaptive group also reported higher usability (SUS: 85.74 ± 7.87 vs. 70.67 ± 10.29) and completed more sessions (11.45 ± 0.50 vs. 9.62 ± 1.08 ; both $p < 0.001$). The results demonstrate the clinical and human-factors value of closed-loop, patient-centered VR adaptation for anxiety and depression treatment.

Keywords—virtual reality therapy; deep Q-learning; patient-centered design; adaptive user interface; anxiety and depression.

1. INTRODUCTION

Anxiety and depressive disorders impose a substantial clinical and socioeconomic burden [1]. Although pharmacotherapy and cognitive-behavioral therapy (CBT) are evidence-based treatments, treatment delivery can be constrained by adherence, dropout, and the limited capacity of fixed protocols to accommodate variability in patients' symptoms, preferences, and readiness [2]. Engineering therapeutic interfaces that respond to measured patient states is therefore clinically relevant.

Virtual reality (VR) provides a controllable, immersive channel for delivering exposure-based and cognitive-behavioral interventions. Recent syntheses indicate that VR-supported therapy can reduce anxiety symptoms and can be comparable with conventional exposure or CBT when implemented in appropriate clinical settings [3]. However, many VR interventions retain a predetermined scenario

progression. Fixed content can under-challenge some users while producing excessive distress or disengagement in others, particularly when arousal fluctuates during treatment.

Real-time physiological and behavioral sensing can extend VR from a predefined presentation medium to a closed-loop therapeutic interface. Deep learning provides representations for high-dimensional multimodal data [4], while adaptive-UI work has shown the feasibility of combining deep Q-learning (DQL) with Golden Jackal Optimization (GJO) for personalized interface control [5]. Jiang et al. reported 45-epoch convergence for their DQL-GJO wearable-interface model, compared with 70 epochs for standard DQL, together with a task-completion time of 82 s, an error rate of 9.9%, and user satisfaction of 78% [5]. Therapeutic VR additionally requires that interface actions conform to exposure-based treatment principles.

Existing systems rarely integrate multimodal sensing, reinforcement-based interface control, and patient-centered VR treatment in a single clinical workflow. In particular, an adaptive policy must regulate stimulation so that exposure remains therapeutically challenging yet tolerable, rather than merely minimizing arousal. The engineering problem is to translate multimodal observations into safe, interpretable environmental actions that can be evaluated against a fixed-progression VR control.

This study designs and evaluates a closed-loop adaptive VR therapeutic system for adults with generalized anxiety disorder (GAD) or major depressive disorder (MDD). The system (1) fuses EEG, heart rate, galvanic skin response, gaze, and movement signals into a state representation; (2) uses a CNN-LSTM DQL model, with GJO-assisted hyperparameter initialization, to select discrete environmental adjustments; and (3) evaluates clinical symptoms, usability, and session completion in a 12-week randomized controlled trial (RCT). The contribution is a patient-centered therapeutic-interface architecture whose adaptation target is a clinically tolerable exposure range rather than a static scenario schedule.

*Malambo Lushomo, Zambia Precision Engineering Limited Liability Company, Lusaka, Zambia, malamblo@outlook.com

2. RELATED WORK

VR psychotherapy has advanced from static simulations to immersive, interactive interventions, while AI-enabled interfaces offer mechanisms for individualized control. This section compares evidence on VRET, adaptive interaction, and patient-centered digital mental health, then identifies the missing integration.

2.1. *Virtual Reality in Psychotherapy*

Virtual reality exposure therapy (VRET) creates controllable, repeatable exposure contexts and is commonly used for anxiety-related conditions. Systematic reviews have reported clinically meaningful reductions in anxiety, while the comparative value of VRET relative to in-vivo exposure varies by indication, comparator, and study design [3], [6]–[8]. Modern head-mounted displays also make scenario intensity, timing, and sensory cues adjustable within clinical sessions.

Recent evidence has expanded the focus beyond symptom change to safety, usability, acceptability, and attrition. Systematic reviews and clinical trials report generally favorable effectiveness and user-experience outcomes, while identifying simulator sickness, exposure-related distress, and inconsistent user-experience measurement as design concerns [9]–[13]. These findings support VR as a configurable treatment delivery platform and underscore the need for careful management of safety and user-experience outcomes.

2.2. *Deep Learning in Medical Device Interaction*

Deep learning and reinforcement learning enable adaptive interfaces that map observed states to actions. For wearable medical devices, DQL-GJO has been used to adjust navigation flow, button placement, and notification timing from real-time user behavior [5]. In therapeutic settings, reinforcement learning has been investigated for dynamic treatment regimes and precision medicine, indicating that the state–action–reward formulation can support sequential clinical decisions when objectives and constraints are explicitly defined [14]. The broader machine-learning literature also establishes the value of deep representations for high-dimensional data and adaptive interaction [4], [15]–[18].

Therapeutic VR differs from ordinary UI personalization because actions change the intensity and structure of clinically meaningful stimuli. Therefore, the adaptation policy must incorporate a safety-aware reward definition, use a limited discrete action set, and preserve a graded exposure trajectory. Existing VR literature focuses mainly on fixed scenario delivery and clinical outcomes [3], [6]–[13], whereas general adaptive-UI work does not address exposure pacing, multimodal affective input, or clinical endpoints [5], [14], [16], [17].

2.3. *Patient-Centered Interaction Design*

Patient-centered design requires systems to respond to patients' functional needs, preferences, and changing capacity during care. In digital mental health, perceived usefulness, usability, and sustained engagement are central implementation outcomes [9], [19]–[23]. For immersive interventions, an interface must balance agency and guidance: it should adjust task difficulty without obscuring treatment goals or removing clinician oversight.

Physiological and behavioral signals can be used as proxies for momentary interaction states, but they must be interpreted within the treatment context. VR exposure relies

on systematic engagement with feared cues and subsequent learning; the appropriate target is therefore a tolerable therapeutic arousal range rather than minimal arousal. Affective-computing, emotion-recognition, and therapeutic-alliance studies support multimodal and patient-centered interaction design, while contemporary VR reviews emphasize safety and standardized user-experience assessment [9], [10], [24]–[28].

2.4. *Research Gap and Contribution*

Three gaps emerge. First, many clinical VR interventions use fixed progression rather than state-aware adaptation [3], [6]–[13]. Second, adaptive-interface studies show the technical feasibility of DQL-based personalization but usually evaluate navigation or task-performance outcomes rather than therapeutic exposure [5], [14]. Third, the link between multimodal sensing, exposure pacing, and patient-centered clinical evaluation remains underdeveloped [29]. An integrated system must align data-driven control with a clinically tolerable arousal range, measurable user experience, and symptom outcomes.

This study addresses these gaps with a closed-loop patient-centered VR platform. Its technical contributions are a multimodal state representation, a CNN–LSTM DQL policy with GJO-assisted hyperparameter initialization, and a safety-oriented reward formulation for discrete environmental adaptation. Its clinical contribution is a randomized comparison of adaptive and fixed-progression VR using symptoms, usability, and treatment completion as outcomes. The platform is implemented as a clinician-supervised therapeutic-interface framework.

3. METHODS

We conducted a two-arm, parallel-group RCT to evaluate a deep learning–optimized adaptive VR interface against an otherwise identical fixed-progression VR condition. The protocol was approved by the institutional review board, and all participants provided written informed consent. Clinical assessments were conducted at baseline and at scheduled follow-up visits in weeks 4, 8, and 12.

3.1. *Participants*

A total of 120 participants were recruited from outpatient psychiatric clinics. Inclusion criteria were: (1) age 18–65 years; (2) a primary diagnosis of GAD or MDD according to DSM-5 criteria; and (3) a baseline score of at least 14 on the Hamilton Anxiety Rating Scale (HAMA) or Hamilton Depression Rating Scale (HAMD). Exclusion criteria were: (1) a history of psychosis or bipolar disorder; (2) current substance-use disorder; (3) significant suicide risk; and (4) photosensitive epilepsy or severe motion sickness. All participants provided written informed consent before enrollment.

Eligible participants were randomly allocated to the Adaptive VRT group (n=60) or the Non-Adaptive VRT group (n=60) using a computer-generated, diagnosis-stratified sequence (GAD or MDD). Allocation was completed before the first intervention session, and each participant followed the same schedule of clinical assessments.

3.2. *System Architecture*

The therapeutic platform comprises sensing, processing, adaptive-control, and VR-delivery modules.

Hardware. A Meta Quest 3 headset provided visual and auditory immersion. A non-invasive EEG headband acquired alpha, beta, theta, and delta frequency-band features, and a wrist-worn sensor acquired heart rate (HR) and galvanic skin response (GSR). Sensor streams were synchronized with in-VR event timestamps.

Software. Unity-based therapeutic environments included low-stimulation nature scenes (e.g., forest and beach) and graded social-exposure scenes (e.g., crowd settings and public speaking). The Non-Adaptive VRT condition used a predefined scenario schedule. The Adaptive VRT condition used the same scenario library, but environmental parameters were controlled by the adaptive engine.

Closed-loop integration. As shown in Figure 1, the platform streams physiological sensing and in-VR interaction data, including gaze direction and movement speed, to the processing layer. The processor constructs a state vector and submits it to the DQL policy. The selected action adjusts VR parameters and is logged with the updated state, producing a closed-loop sequence of observation, decision, rendering, and response.




Figure 1. Closed-loop architecture of the adaptive VR therapeutic platform.

Physiological and behavioral data are continuously acquired and aligned in time. The state vector is passed to the DQL policy, which selects one discrete adaptation from a constrained action set. The rendered parameters include lighting, crowd density, ambient sound, and interactive elements. The constrained loop preserves the therapeutic content library while adapting pacing and stimulus intensity to the participant's observed state.

3.3. The Adaptive Interface Model

We formulated interface adaptation as a finite-action reinforcement-learning problem. At time t , the policy observes state s_t , selects action a_t , receives reward r_t , and updates the Q-value estimate for subsequent decisions [18].

3.3.1. DQL Formulation

State (S). At decision time t , s_t combines the preceding 30-s window of physiological features (EEG band features, HR, and GSR) and behavioral features (gaze, head-movement velocity, and interaction frequency). The state representation preserves the temporal context needed to select an environmental adjustment.

Action (A). The action set contains discrete, bounded adjustments to lighting and weather, virtual-crowd density, ambient-sound volume and type, and interactive elements. Each action changes the current scene configuration without introducing an unplanned therapeutic scenario.

Reward (R). The reward function is designed to maintain exposure within a clinically tolerable arousal range. Positive reward is assigned when the participant remains engaged and physiological arousal remains within a predefined target interval during the planned exposure step. Reduced reward is assigned to abrupt arousal escalation, sustained disengagement, or an intensity change that moves the participant outside the target interval. Thus, the policy does not reward indiscriminate arousal reduction; it rewards safe engagement with graded exposure.

Network architecture. Figure 2 shows a hybrid DQL architecture. A 1D CNN extracts local multivariate patterns from physiological data (EEG, HR, and GSR) over the observation window. An LSTM models temporal dependencies in the behavioral sequence (gaze and movement). The branch embeddings are concatenated and passed through fully connected layers to approximate $Q(s_t, a)$ for each permissible action.

The CNN and LSTM branches retain complementary signal characteristics: the CNN encodes local co-variation in the multivariate physiological window, whereas the LSTM retains temporal structure in gaze and movement trajectories. The fused representation is used to estimate Q-values across the discrete action space, allowing the controller to choose the highest-valued permissible adaptation.

GJO-assisted optimization. GJO was used during model configuration to search the hyperparameter space and establish the exploration-exploitation schedule for the DQL agent [30]. The clinical trial evaluates the integrated adaptive system, which combines these components.




Figure 2. Hybrid CNN-LSTM DQL architecture for adaptive VR interface control.

3.4. Procedure

All participants completed a 12-week treatment program consisting of one 45-min VR session per week. The Adaptive VRT group received real-time, constrained environmental adjustments from the DQL-based engine. The Non-Adaptive VRT group used identical hardware, session duration, and scenario library, but scenarios followed a fixed schedule without real-time adaptation. This design isolates real-time adaptation within the shared VR program.

3.5. Outcome Measures

Primary outcomes were changes in anxiety and depression symptoms from baseline to week 12, measured by HAMA and the 17-item HAMD. The repeated-measures model also included scheduled assessments at weeks 4 and 8.

Secondary outcomes included: (1) the Patient Health Questionnaire-9 (PHQ-9) and Generalized Anxiety Disorder-7 (GAD-7) for self-reported symptoms; (2) the System Usability Scale (SUS) for perceived usability; (3) treatment adherence, measured as completed sessions; and (4) system-use data, including session duration and interaction frequency.

3.6. Statistical Analysis

Data were analyzed using SPSS version 28.0. Baseline group comparisons used independent-samples *t*-tests. Repeated HAMA and HAMD measures were analyzed with mixed-effects models that included fixed effects for group, time (baseline, week 4, week 8, and week 12), and their interaction, with participant-level repeated observations modeled as a random effect. Planned post hoc group comparisons were evaluated with Bonferroni adjustment within each outcome family; all tests used two-sided $\alpha=0.05$. SUS and completed sessions were compared between groups using independent-samples *t*-tests. The primary model is

reported as a group $\times$ time interaction, and week-12 comparisons are reported as endpoint analyses.

4. RESULTS

4.1. Participant Characteristics

Baseline clinical characteristics are summarized in Table 1. The two groups each contained 60 participants. Baseline HAMA and HAMD scores did not differ significantly between groups (both $p>0.05$), supporting comparability for the clinical severity measures reported in Table 1.

TABLE 1. BASELINE CLINICAL CHARACTERISTICS OF PARTICIPANTS

Group	Baseline HAMA (Mean $\pm$ SD)	Baseline HAMD (Mean $\pm$ SD)
Adaptive	24.23 $\pm$ 4.54	24.46 $\pm$ 4.98
Non-Adaptive	25.02 $\pm$ 4.63	24.23 $\pm$ 4.99

4.2. Primary Outcomes: Clinical Efficacy

Both groups showed symptom reductions over the 12-week treatment period; however, the Adaptive VRT group showed a larger reduction than the Non-Adaptive VRT group.

Mixed-effects analysis showed a significant group $\times$ time interaction for HAMA, $F(3, 354)=8.82$, $p<0.001$, and HAMD, $F(3, 354)=7.15$, $p<0.001$. Figure 3 presents the group trajectories across the scheduled assessments.

At week 12, mean HAMA was 12.03 (SD=5.65) in the Adaptive VRT group and 17.93 (SD=5.42) in the Non-Adaptive VRT group, $t(118)=-5.49$, $p<0.001$. Mean HAMD was 12.45 (SD=5.85) in the Adaptive VRT group and 15.98 (SD=5.33) in the Non-Adaptive VRT group, $t(118)=-3.21$, $p=0.002$.




Figure 3. HAMA and HAMD trajectories for adaptive and non-adaptive VRT over 12 weeks.

(A) Lines connect group means at baseline and scheduled assessments (weeks 4, 8, and 12); shaded bands represent standard deviations. The HAMA group $\times$ time interaction was significant, $F(3, 354)=8.82$, $p<0.001$. (B) The corresponding HAMD group $\times$ time interaction was significant, $F(3, 354)=7.15$, $p<0.001$. Week-12 group differences are reported in the main text.

4.3. Secondary Outcomes: Usability and Adherence

At week 12, SUS was higher in the Adaptive VRT group (mean=85.74, SD=7.87) than in the Non-Adaptive VRT group (mean=70.67, SD=10.29), $t(118)=8.51$, $p<0.001$ (Figure. 4).




Figure 4. Distribution of System Usability Scale (SUS) scores by intervention group.

The box plots show the median, interquartile range, whiskers, and individual outlier for each group. Higher SUS scores indicate better perceived usability and satisfaction; the Adaptive VRT group had a higher group mean than the Non-Adaptive VRT group.

Treatment adherence was also higher in the Adaptive VRT group: participants completed an average of 11.45 of 12 sessions ($SD=0.50$), compared with 9.62 sessions ($SD=1.08$) in the Non-Adaptive VRT group, $t(118)=10.8$, $p<0.001$. Table 2 summarizes the week-12 outcomes.

TABLE II. WEEK-12 CLINICAL, USABILITY, AND ADHERENCE OUTCOMES

Group	Week 12 HAMA (Mean $\pm$ SD)	Week 12 HAMD (Mean $\pm$ SD)	SUS Score (Mean $\pm$ SD)	Adherence (Sessions, Mean $\pm$ SD)
Adaptive	12.03 $\pm$ 5.65	12.45 $\pm$ 5.85	85.74 $\pm$ 7.87	11.45 $\pm$ 0.50
Non-Adaptive	17.93 $\pm$ 5.42	15.98 $\pm$ 5.33	70.67 $\pm$ 10.29	9.62 $\pm$ 1.08

5. DISCUSSION

The adaptive interface was associated with lower week-12 HAMA and HAMD scores, higher SUS, and more completed sessions than fixed-progression VRT. The findings show that aligning scenario delivery with real-time sensing can improve the clinical and human-factors outcomes measured in this trial.

5.1. Interpretation of Key Findings

The significant group $\times$ time interactions show that symptom trajectories differed between Adaptive and Non-Adaptive VRT. Both groups used the same VR content library and session length; the key difference was whether environmental parameters responded to the participant's observed state. The pattern is consistent with the clinical utility of individualized pacing during immersive intervention. In exposure-based work, the goal is not to eliminate arousal but to support safe, repeated engagement with feared stimuli at a tolerable intensity. The reward formulation operationalizes this principle by favoring engagement within a target arousal range while penalizing abrupt escalation or disengagement. The result links patient-centered interaction control with graded exposure pacing.

The higher SUS scores and greater session completion in the Adaptive VRT group indicate that usability and adherence

can improve concurrently when the interface responds to measured participant states. This finding is important for digital mental-health interventions, in which sustained engagement is an implementation requirement as well as a user-experience outcome [9], [19], [21]–[23].

At the technical level, the hybrid CNN–LSTM controller integrates physiological and behavioral channels into the state representation, and GJO supports configuration of the hyperparameter and exploration–exploitation settings. This design is motivated by the DQL-GJO adaptive-interface framework reported for wearable medical devices [5] and by broader reinforcement-learning applications in healthcare [14]. The integrated controller serves as the intervention mechanism in the clinical evaluation, with algorithm-level benchmarking as a complementary engineering evaluation.

5.2. Clinical and Theoretical Implications

Clinically, a closed-loop VR interface can support standardized delivery of graded stimuli while retaining individualized pacing. In outpatient settings, this can reduce the need to manually adjust every environmental parameter and provide a structured record of state–action trajectories for clinician review. The platform keeps VR therapy within a clinician-supervised workflow and treats usability as a design variable rather than a secondary implementation issue.

The work illustrates a computational patient-centered design approach: sensing provides a state estimate, a reinforcement-learning policy selects a bounded action, and clinical outcomes determine the value of the therapeutic interface. This formulation couples human–computer interaction metrics with clinical endpoints and establishes an engineering basis for evaluating adaptive immersive interventions [28].

5.3. Limitations and Future Research

Several limitations should be considered. The sample was recruited from outpatient clinics and may not represent the broader population with anxiety and depression. The 12-week intervention established short-term outcomes but did not assess the durability of symptom changes. Multicentre trials with longer follow-up are needed to evaluate generalizability and maintenance effects.

The clinical comparison evaluated an integrated adaptive system against fixed-progression VRT. Future engineering evaluations should include pre-specified baselines, such as rule-based adaptation, DQL without GJO, and DQL-GJO, together with learning curves, cumulative-reward trajectories, and component ablations for the CNN and LSTM branches. These analyses will isolate the contributions of the controller, signal modalities, and optimization strategy. Future trials should also include an active CBT comparator and longer follow-up.

The state representation can be extended with additional clinically appropriate signals, including structured clinician ratings, speech features, or facial-expression features. Future work should examine how adaptation affects engagement, therapeutic arousal, self-efficacy, and symptom change across anxiety, depression, post-traumatic stress disorder, phobias, and substance-use conditions.

6. CONCLUSION

This study developed a patient-centered, closed-loop VR therapeutic system that uses multimodal sensing and CNN–

LSTM DQL control to adapt virtual environments during treatment. In 120 participants, Adaptive VRT was associated with lower week-12 HAMA and HAM-D scores, higher usability, and greater session completion than fixed-progression VRT. The reward design aligns environmental adaptation with a tolerable therapeutic arousal range, enabling graded exposure rather than simple arousal suppression. The findings support the engineering value of integrating state-aware interface control with clinical and usability evaluation in immersive mental-health interventions.

REFERENCES

- [1] World Health Organization, "Mental health: Strengthening our response," 2022. [Online]. Available: <https://www.who.int/news-room/fact-sheets/detail/mental-health-strengthening-our-response>
- [2] A. E. Kazdin, "Addressing the treatment gap: A key challenge for extending evidence-based psychosocial interventions," *Behav. Res. Ther.*, vol. 88, pp. 7–18, 2017, doi: 10.1016/j.brat.2016.06.004.
- [3] W. Zeng, J. Xu, J. Yu, and X. Chu, "Effectiveness of virtual reality therapy in the treatment of anxiety disorders in adolescents and adults: A systematic review and meta-analysis of randomized controlled trials," *Front. Psychiatry*, vol. 16, Art. no. 1553290, 2025, doi: 10.3389/fpsy.2025.1553290.
- [4] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015, doi: 10.1038/nature14539.
- [5] M. Jiang, J. Huang, and L. Wang, "Adaptive user interfaces for wearable medical devices using deep Q-learning and Golden Jackal Optimization," *Sci. Rep.*, vol. 15, Art. no. 44776, 2025, doi: 10.1038/s41598-025-28937-z.
- [6] J. Wu, Y. Sun, G. Zhang, Z. Zhou, and Z. Ren, "Virtual reality-assisted cognitive behavioral therapy for anxiety disorders: A systematic review and meta-analysis," *Front. Psychiatry*, vol. 12, Art. no. 575094, 2021, doi: 10.3389/fpsy.2021.575094.
- [7] T. Horigome, S. Kurokawa, K. Sawada, S. Kudo, K. Shiga, M. Mimura, and T. Kishimoto, "Virtual reality exposure therapy for social anxiety disorder: A systematic review and meta-analysis," *Psychol. Med.*, vol. 50, no. 15, pp. 2487–2497, 2020, doi: 10.1017/S0033291720003785.
- [8] N. Baghaei, V. Chitale, A. Hlasnik, L. Stemmet, H.-N. Liang, and R. Porter, "Virtual reality for supporting the treatment of depression and anxiety: Scoping review," *JMIR Ment. Health*, vol. 8, no. 9, e29681, 2021, doi: 10.2196/29681.
- [9] S. Shahid, J. Kelson, and A. Saliba, "Effectiveness and user experience of virtual reality for social anxiety disorder: Systematic review," *JMIR Ment. Health*, vol. 11, e48916, 2024, doi: 10.2196/48916.
- [10] P. Wang, X. Ai, X. Zhang, F. Ma, Y. Zhuang, and S. Wang, "Evaluating virtual reality technology in psychotherapy: Impacts on anxiety, depression, and ADHD," *Front. Psychiatry*, vol. 15, Art. no. 1480788, 2024, doi: 10.3389/fpsy.2024.1480788.
- [11] M. Lee et al., "Virtual reality-based cognitive behavior therapy for major depressive disorder: An alternative to pharmacotherapy for reducing suicidality," *Yonsei Med. J.*, vol. 66, no. 1, pp. 25–36, 2025, doi: 10.3349/ymj.2024.0002.
- [12] Y. Pu and H. Luo, "Using virtual reality technology to treat adolescent and young adult depression disorder: A systematic review and meta-analysis," *Discover Psychol.*, vol. 5, no. 1, Art. no. 58, 2025, doi: 10.1007/s44202-025-00384-x.
- [13] D. Freeman et al., "Automated virtual reality therapy to treat agoraphobic avoidance and distress in patients with psychosis (gameChange): A multicentre, parallel-group, single-blind, randomised, controlled trial in England with mediation and moderation analyses," *Lancet Psychiatry*, vol. 9, no. 5, pp. 375–388, 2022, doi: 10.1016/S2215-0366(22)00060-8.
- [14] T. C. Frommeyer et al., "Reinforcement learning and its clinical applications within healthcare: A systematic review of precision medicine and dynamic treatment regimes," *Healthcare*, vol. 13, no. 14, Art. no. 1752, 2025, doi: 10.3390/healthcare13141752.
- [15] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Netw.*, vol. 61, pp. 85–117, 2015, doi: 10.1016/j.neunet.2014.09.003.
- [16] K. Z. Gajos, J. O. Wobbrock, and D. S. Weld, "Improving the performance of motor-impaired users with automatically-generated, ability-based interfaces," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, 2008, pp. 1257–1266, doi: 10.1145/1357054.1357250.
- [17] S. Amershi, M. Cakmak, W. B. Knox, and T. Kulesza, "Power to the people: The role of humans in interactive machine learning," *AI Mag.*, vol. 35, no. 4, pp. 105–120, 2014, doi: 10.1609/aimag.v35i4.2513.
- [18] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Mach. Learn.*, vol. 8, no. 3, pp. 279–292, 1992, doi: 10.1007/BF00992698.
- [19] H. L. O'Brien and E. G. Toms, "What is user engagement? A conceptual framework for defining user engagement with technology," *J. Amer. Soc. Inf. Sci. Technol.*, vol. 59, no. 6, pp. 938–955, 2008, doi: 10.1002/asi.20801.
- [20] M. Slater and S. Wilbur, "A framework for immersive virtual environments (FIVE): Speculations on the role of presence in virtual environments," *Presence*, vol. 6, no. 6, pp. 603–616, 1997, doi: 10.1162/pres.1997.6.6.603.
- [21] S. Barelllo, G. Graffigna, and E. Vegni, "Patient engagement as an emerging challenge for healthcare services: Mapping the literature," *Nurs. Res. Pract.*, vol. 2012, Art. no. 905934, 2012, doi: 10.1155/2012/905934.
- [22] G. Eysenbach, "What is e-health?" *J. Med. Internet Res.*, vol. 3, no. 2, e20, 2001, doi: 10.2196/jmir.3.2.e20.
- [23] S. R. Steinhubl, E. D. Muse, and E. J. Topol, "Can mobile health technologies transform health care?" *JAMA*, vol. 310, no. 22, pp. 2395–2396, 2013, doi: 10.1001/jama.2013.281078.
- [24] R. W. Picard, *Affective Computing*. Cambridge, MA, USA: MIT Press, 2000.
- [25] X. Chen, M. Cao, H. Wei, Z. Shang, and L. Zhang, "Patient emotion recognition in human–computer interaction system based on machine learning method and interactive design theory," *J. Med. Imaging Health Inform.*, vol. 11, no. 2, pp. 307–312, 2021, doi: 10.1166/jmhi.2021.3293.
- [26] E. S. Bordin, "The generalizability of the psychoanalytic concept of the working alliance," *Psychother.: Theory, Res. Pract.*, vol. 16, no. 3, p. 252, 1979, doi: 10.1037/h0085885.
- [27] A. Miloff, P. Carlbring, W. Hamilton, G. Andersson, L. Reuterskiöld, and P. Lindner, "Measuring alliance toward embodied virtual therapists in the era of automated treatments with the Virtual Therapist Alliance Scale (VTAS): Development and psychometric evaluation," *J. Med. Internet Res.*, vol. 22, no. 3, e16660, 2020, doi: 10.2196/16660.
- [28] P. R. Montague, R. J. Dolan, K. J. Friston, and P. Dayan, "Computational psychiatry," *Trends Cogn. Sci.*, vol. 16, no. 1, pp. 72–80, 2012, doi: 10.1016/j.tics.2011.11.018.
- [29] Z. Azizi et al., "Challenge of optimizing medical therapy in heart failure: Unlocking the potential of digital health and patient engagement," *J. Amer. Heart Assoc.*, vol. 13, no. 2, e030952, 2024, doi: 10.1161/JAHA.123.030952.
- [30] N. Chopra and M. M. Ansari, "Golden jackal optimization: A novel nature-inspired optimizer for engineering applications," *Expert Syst. Appl.*, vol. 198, Art. no. 116924, 2022, doi: 10.1016/j.eswa.2022.116924.

ACKNOWLEDGEMENTS

None.

FUNDING

None.

AVAILABILITY OF DATA

Not applicable.

ETHICAL STATEMENT

All participants provided written informed consent prior to participation. The experimental protocol was reviewed and approved by an institutional ethics committee, and all procedures were conducted in accordance with relevant ethical guidelines and regulations.

AUTHOR CONTRIBUTIONS

Malambo Lushomo conceived and designed the study, developed the patient-centered adaptive virtual reality therapeutic framework, designed the closed-loop VR system architecture, and supervised the integration of multimodal sensing with the deep Q-learning optimization strategy, while conducting the experimental implementation, clinical evaluation, data analysis, and manuscript preparation.

COMPETING INTERESTS

The authors declare no competing interests.

Publisher's note WEDO remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is published online with Open Access by BIG.D and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).

© The Author(s) 2026