applied sciences

MDPI

Perspective

AI-Enhanced Extended Reality for Rehabilitation in Africa: A Perspective on Explainable Agents, Co-Creation, and Generative Worlds

Chala Diriba Kenea [1] and Bruno Bonnechère [2,3,4,*]

[1] Department of Information Science, Faculty of Computing and Informatics, Jimma Institute of Technology, Jimma University, Jimma P.O. Box 378, Ethiopia; chala.diriba@ju.edu.et

[2] REVAL Rehabilitation Research Center, Faculty of Rehabilitation Sciences, Hasselt University, 3590 Diepenbeek, Belgium

[3] Technology-Supported and Data-Driven Rehabilitation, Data Sciences Institute, Hasselt University, 3590 Diepenbeek, Belgium

[4] Department of PXL—Healthcare, PXL University of Applied Sciences and Arts, 3500 Hasselt, Belgium

* Correspondence: bruno.bonnechere@uhasselt.be

Abstract

The burden of disability is rising rapidly in Africa, where a severe shortage of rehabilitation professionals and limited infrastructure create a major treatment gap. Immersive virtual reality and serious games have shown promise for upper limb rehabilitation, but current extended reality (XR) solutions lack personalization, cultural adaptability, real-time feedback, and scalability. This perspective paper proposes a conceptual AI-enhanced XR framework tailored to African low- and middle-income countries. We identify how generative AI, large language models, multiagent systems, and explainable AI can address specific rehabilitation barriers. The framework integrates these four pillars into a three-layer architecture covering content creation, interaction, and decision support. We analyze implementation considerations specific to African contexts—infrastructure, capacity building, cultural adaptation, ethics, and financing—and outline a detailed research agenda with near, medium, and longer term priorities. Realizing this vision requires co-design with African communities, investment in local capacity, adaptation to infrastructure constraints, and development of ethical frameworks. AI-enhanced XR has the potential to democratize access to quality rehabilitation across Africa, but this potential must be validated through rigorous, context-sensitive research.

Keywords: artificial intelligence; extended reality; rehabilitation; low- and middle-income countries; Africa; generative AI; explainable AI; multiagent systems

Academic Editors: Marina Anca Cidota and Clara Maathuis

Received: 1 April 2026
Revised: 13 May 2026
Accepted: 13 May 2026
Published: 15 May 2026

1. Introduction

Stroke and other disabling conditions impose a staggering health burden in low- and middle-income countries (LMICs), where two-thirds of stroke-related deaths occur and disability-adjusted life years are seven times higher than in high-income countries [1,2]. Sub-Saharan Africa (SSA) faces a particularly acute shortage of rehabilitation professionals, with ratios as low as 0.5 therapists per 10,000 population in some regions [3]. This scarcity, compounded by limited infrastructure and geographical barriers, leaves the majority of stroke survivors without adequate upper limb rehabilitation—a critical factor for regaining independence and quality of life. The World Health Organization estimates that over 2.4 billion people globally need rehabilitation services, with the largest gaps in LMICs,

 https://doi.org/10.3390/app16104946

and this need is expected to grow due to population aging and the rising prevalence of non-communicable diseases [4–6].

Digital health technologies, especially extended reality (XR), an umbrella term encompassing virtual reality (VR), augmented reality (AR), and mixed reality (MR), represents a continuum of technologies that blend the physical and digital worlds to varying degrees [7,8]. While VR immerses users in a completely synthetic environment, AR overlays digital information onto the real world, and MR enables interactive digital objects to coexist and interact with the physical environment. They offer engaging, repetitive, task-oriented training with real-time feedback and can be deployed in home or community settings, thereby extending the reach of rehabilitation services [9,10]. A growing body of evidence supports the efficacy of VR-based interventions for upper limb recovery after stroke, with benefits including improved range of motion, motor control, and activities of daily living [11,12]. Moreover, serious games have been shown to increase patient motivation and adherence, which are key determinants of rehabilitation success [13].

An example of culturally adapted XR rehabilitation is the AdaptRehab VR system, developed through a participatory co-creation approach, combining patients, therapists and IT specialists in Ethiopia [14]. This system includes six imVR games targeting different upper limb functions (e.g., Basket Bloom for reaching and grasping, Bean Picker Pro for fine motor skills) and was designed in the local Afaan Oromoo language using culturally familiar objects such as coffee beans.

Other LMIC-focused XR rehabilitation initiatives have emerged in recent years. In India, a low-cost VR system for stroke rehabilitation demonstrated feasibility in community settings [15]. Studies from Brazil and China have similarly emphasized the importance of cultural adaptation and the need for AI-enhanced personalization [16,17]. A systematic review of serious games for rehabilitation in LMICs identified 23 studies, concluding that while feasibility has been demonstrated, scalability and personalization remain key barriers [18].

Yet, despite these advances, current XR rehabilitation systems face inherent limitations. They often lack deep personalization to individual patient progress, struggle to scale content across diverse cultural and linguistic contexts, provide only rudimentary feedback, and do not exploit the full potential of data-driven adaptation [19,20]. These limitations become even more critical in LMICs where the therapist-to-patient ratio is extremely low, and where the need for autonomous, intelligent, and culturally responsive systems is paramount [21,22].

Recent breakthroughs in artificial intelligence (AI) offer a paradigm shift. Generative AI can produce infinite variations of therapeutic environments and tasks, tailored to a patient's culture and functional level [23]. Large language models (LLMs) enable natural, empathetic dialog that can simulate a therapist's guidance [24]. Multiagent systems (MAS) can introduce virtual peers or coaches that scale social interaction and motivation [25]. Explainable AI (XAI) can build trust by making the system's reasoning transparent to both patients and clinicians [26]. When integrated with XR, these AI capabilities promise to transform rehabilitation from a one-size-fits-all, therapist-intensive model into a personalized, scalable, and equitable service—especially in underserved regions.

This perspective paper argues that AI-enhanced XR can help bridging the rehabilitation gap in Africa. We first summarize the current state of XR rehabilitation in LMICs. Then we explore in depth how each of these four AI techniques—generative AI, LLMs, MAS, and XAI—can address specific barriers to access, personalization, and trust. We present a conceptual framework for an AI-XR rehabilitation ecosystem tailored to African contexts, discuss implementation challenges and ethical considerations, and propose a detailed research agenda to turn this vision into reality.

2. The Current Landscape: XR Rehabilitation in LMICs

The potential of serious games and imVR for rehabilitation in LMICs has been increasingly recognized. A recent bibliometric analysis revealed that while high-income countries dominate research output, a growing number of studies from Africa, the Middle East, and South Asia are emerging [18]. However, the vast majority of existing XR solutions are developed in high-income settings and are not designed with the cultural, linguistic, or infrastructural realities of LMICs in mind [27]. This mismatch can lead to low acceptance, poor engagement, and limited effectiveness when these technologies are imported without adaptation [28].

Together, these studies converge on several key findings: (i) XR rehabilitation is feasible and acceptable in LMIC settings when culturally adapted; (ii) infrastructure challenges (power, internet) are significant but not insurmountable; (iii) therapist time remains a bottleneck, limiting scalability; (iv) personalization and content adaptation are manual and labor-intensive; and (v) none of these systems incorporate advanced AI for dynamic adaptation.

The AdaptRehab VR project, conducted in Ethiopia, represents a deliberate departure from this trend. Through a human-centered design approach involving patients, physiotherapists, and local stakeholders, six imVR games were co-created, each targeting specific upper limb functions (e.g., range of motion, coordination, fine motor skills). The system was developed in the local Afaan Oromoo language, used culturally familiar objects (e.g., coffee beans), and employed both controllers and hand tracking to accommodate varying levels of impairment. A Technology Acceptance Model (TAM) evaluation with 10 participants (4 clinicians, 6 patients) yielded high scores for perceived ease of use (clinicians 4.5/5, patients 4.2/5) and perceived usefulness (clinicians 4.7/5, patients 4.3/5), demonstrating initial acceptance in this small sample size study [14]. This co-creation approach aligns with broader recommendations for developing culturally relevant digital health interventions in LMICs [29,30].

Yet, we identified important limitations: the system currently offers only a limited number of games, customization is primarily manual (though with automatic progression), and the linguistic adaptation is confined to a single Ethiopian language. Moreover, feedback is limited to simple audio-visual cues and score increments, and the system does not provide the kind of natural, conversational guidance that a human therapist would offer [14]. These limitations reflect a broader challenge: while XR can deliver therapy exercises, it cannot yet replicate the adaptive, reasoning, and empathic qualities of a skilled clinician—qualities that are in desperately short supply in African LMICs [3,31].

3. AI as a Catalyst for Accessible Rehabilitation in Africa

The convergence of XR with advanced AI techniques offers a pathway to overcome the current limitations and unlock scalable, personalized rehabilitation. Below, we explore in detail how each of the four AI pillars might be harnessed to address specific challenges in African LMICs, with a focus on reducing dependency on scarce human resources and enhancing cultural and linguistic appropriateness. For each technology, we present both opportunities and critical limitations, explicitly acknowledging where evidence is currently lacking.

3.1. Generative AI for Culturally Adaptive Content Creation

One of the most significant barriers to deploying XR rehabilitation in Africa is the cost and effort required to create culturally relevant content. Manual development of 3D assets, environments, and audio instructions is labor-intensive and requires specialized skills that are often not available locally [32]. As a result, most XR applications imported from high-income countries feature Western settings, objects, and languages, which can feel alienating and reduce engagement [28]. Generative AI offers a potential solution by enabling the automated creation of an unlimited variety of culturally appropriate content on demand.

Generative AI models, such as diffusion models for 3D asset generation and text-to-speech systems for voice synthesis, could produce context-specific assets based on simple textual prompts. For example, a therapist might type "create a virtual market scene with mangoes, plantains, and a woven basket" and the AI could generate a fully realized 3D environment with those items [33]. Similarly, voice interfaces could synthesize instructions and encouragement in dozens of African languages, moving beyond the few that can be manually recorded. Recent advances in neural text-to-speech have made it possible to generate natural-sounding voices even for low-resource languages, provided there is sufficient training data [34].

Beyond static assets, generative AI could dynamically create new exercises that target specific movement deficits. We hypothesize that if a patient shows difficulty with shoulder abduction, the system could generate a series of reaching tasks that gradually increase the required range of motion, using objects that are familiar to the patient's daily life—such as reaching for a gourd or a traditional tool [35]. This dynamic content generation ensures that therapy remains challenging and engaging, preventing the monotony that often leads to dropout.

In the African context, where hundreds of languages are spoken and cultural practices vary widely, generative AI can dramatically reduce the development burden. A single platform could serve patients in Ethiopia, Kenya, and Nigeria, each experiencing a version tailored to their own cultural environment. This aligns with the "localized content" strategy identified as critical for engagement in LMICs [18]. Moreover, generative AI can be used to create educational content that explains the rationale behind exercises in culturally appropriate ways, enhancing health literacy—a well-known barrier in many African settings [36].

Critical Considerations

The use of generative AI in clinical contexts raises several concerns.

First, real-time generation of high-quality 3D assets on edge devices (e.g., standalone XR headsets) remains technically challenging due to computational demands; current solutions typically rely on cloud processing, which is problematic in low-connectivity settings [37].

Second, generative models can produce "hallucinated" or unsafe content (e.g., objects that do not exist, anatomically incorrect exercise demonstrations), requiring robust filtering mechanisms [38].

Third, the quality of text-to-3D generation for complex, interactive therapeutic objects is currently far below what is needed for clinical use.

These limitations suggest that near-term implementations will likely use template-based procedural generation augmented by generative AI, rather than fully autonomous content creation.

3.2. Large Language Models (LLMs) for Natural, Empathetic Guidance

The shortage of rehabilitation professionals means that most patients receive little or no personalized coaching. Even when therapists are available, the time they can spend with each patient is limited, and follow-up at home is often nonexistent [3]. Current XR systems provide only simple feedback (e.g., "good job" or a score), lacking the ability to correct movement strategies, explain the rationale behind exercises, or offer motivational support tailored to the patient's emotional state. LLMs, especially those fine-tuned on rehabilitation dialogs, could fill this gap by acting as conversational agents that simulate aspects of a therapist's role.

LLMs such as Chat-GPT, Co-Pilot, Gemini, Claude, Mistral, etc., or open-source alternatives might be adapted to understand patient queries, provide real-time movement corrections based on sensor data, explain how an exercise targets specific functional goals, and adapt their communication style to the patient's mood and literacy level [39]. Because LLMs can handle open-ended dialog, patients could ask questions like "Why does my shoulder hurt when I reach up?" and receive a context-aware explanation that references their recent movement data. This capability could be particularly valuable in settings where patients have low health literacy and may not understand the purpose of exercises [36].

In the African context, LLMs could be fine-tuned with local languages and cultural references to ensure natural communication. While many LLMs are trained primarily on English and other high-resource languages, recent efforts have focused on creating multilingual models and adapting them for low-resource African languages [40]. Through transfer learning and community-based data collection, it may be possible to develop LLM-based coaches that speak Swahili, Yoruba, Amharic, or other widely used languages, making rehabilitation support accessible to patients who are not proficient in English or French.

Furthermore, LLMs can serve as a training tool for community health workers (CHWs) who often supervise patients in rural areas. A CHW could use the LLM to get just-in-time advice on how to assist a patient with a particular exercise, effectively extending the reach of scarce rehabilitation expertise [41]. By providing both direct patient support and capacity building for frontline health workers, LLMs can help bridge the human resource gap.

Critical Considerations

The use of LLMs in clinical rehabilitation presents substantial risks.

First, LLMs are known to "hallucinate", producing confident but incorrect information—which could lead to harmful advice or patient misinformation [39]. In rehabilitation, an incorrect instruction about movement could cause injury.

Second, LLMs fine-tuned on Western data may not understand African cultural norms around health, disability, or communication.

Third, the computational requirements for LLMs are substantial; while quantized versions (e.g., 4-bit Llama 3-8B) can run on high-end smartphones, performance on low-cost XR headsets is unproven. Fourth, the accuracy of LLMs for real-time motion correction has not been validated; this remains a hypothesis requiring empirical testing.

Finally, privacy concerns arise when patient health data is sent to cloud-based LLM services [42].

3.3. Multiagent Systems (MAS) for Scalable Social Support and Collaboration

Rehabilitation can be isolating, and social support is a key determinant of adherence [43]. In LMICs, group therapy is often not available due to resource constraints, and patients may lack peers who understand their condition. Additionally, complex tasks requiring cooperation (e.g., bimanual activities or tasks that involve turn-taking) are difficult to practice alone [44]. Multiagent systems (MAS) can populate the XR environment with intelligent virtual agents (e.g., peers, coaches, or even avatars of family members) that interact with the patient in socially meaningful ways.

These agents could be designed to model social behaviors, provide encouragement, demonstrate exercises, and engage in collaborative tasks. For example, a virtual peer agent could perform the same exercise alongside the patient, offering friendly competition and social comparison, which are known motivators [45]. Agents could also be configured to behave in ways that respect local cultural norms, such as age-based deference or gender roles, making interactions feel natural and acceptable [46]. In cultures where family involvement in care is strong, agents could simulate family members who encourage the patient or explain exercises to visiting relatives.

MAS also could enable the creation of virtual group therapy sessions, even when patients are physically dispersed. Each patient could be represented by an avatar, and AI agents could facilitate the session, ensuring that everyone gets a turn and that the group dynamics remain positive [47]. This is particularly relevant in rural Africa, where traveling to a rehabilitation center may be impossible due to distance and cost.

Moreover, MAS can support telerehabilitation by acting as intermediaries between patients and remote therapists. Agents can summarize patient progress, alert therapists to potential issues, and even coordinate care among multiple providers [48]. By automating routine social interactions and data aggregation, MAS frees up valuable clinician time for more complex tasks, thereby increasing overall system efficiency.

Critical Considerations

MAS for healthcare remains largely experimental.

The emergent behaviors of multiple interacting agents are difficult to predict and control; there are documented cases of simple MAS producing unintended harmful behaviors [49]. In a rehabilitation context, an agent that inadvertently demoralizes a patient or provides incorrect social cues could reduce engagement or cause psychological harm. Additionally, the computational overhead of simulating multiple intelligent agents with natural language capabilities is substantial. Current evidence for MAS in rehabilitation is limited to small laboratory studies; scalability to clinical deployment in LMICs has not been demonstrated.

3.4. Explainable AI (XAI) for Building Trust and Enabling Clinical Oversight

Clinicians and patients may be skeptical of AI-driven decisions if they cannot understand why a particular exercise is recommended, why difficulty has increased, or why progress is being measured in a certain way [26]. Without transparency, trust erodes, and adoption suffers. This is especially critical in LMICs, where there may already be mistrust of technology or where healthcare providers are concerned about being replaced by machines [50]. XAI techniques can provide explanations that accompany system decisions, making the AI's reasoning visible and auditable.

For example, when the system automatically advances a patient to a harder level, it could show a visual comparison of range-of-motion (RoM) data from previous sessions and explain that the improvement meets the threshold for progression. For clinicians, XAI can generate dashboards that highlight which movements are improving and which remain limited, along with confidence estimates and suggested adjustments to the treatment plan [51]. These explanations can be tailored to the user's expertise: a simple visual for a patient, a detailed statistical report for a clinician.

In the African context, where many clinicians may have limited exposure to AI, XAI can serve as an educational tool, gradually building familiarity and trust. By demonstrating that the AI's recommendations are based on clear, verifiable evidence, XAI can help integrate AI-powered systems into clinical workflows without undermining professional autonomy [37]. Additionally, XAI can be used to detect and mitigate bias: if an AI model is found to perform differently across demographic groups, XAI can help identify the source of the bias and guide corrective actions.

https://doi.org/10.3390/app16104946

From a regulatory perspective, XAI is essential for obtaining approval from health authorities, which in many LMICs require a clear understanding of how a medical device arrives at its decisions [52]. By making the AI explainable, the system becomes more accountable and can be more easily integrated into existing healthcare governance structures.

Critical Considerations

XAI is not a panacea.

First, explanations can be "gamed": models can be designed to produce plausible-looking explanations that do not reflect actual reasoning [53].

Second, explanations that are too complex may confuse rather than clarify; those that are too simple may omit critical information.

Third, there is emerging evidence that explanations can create false reassurance: users may trust an AI system more than warranted simply because it provides an explanation [54].

Fourth, different stakeholders (patients, clinicians, regulators) need different types of explanations, and generating all of them adds computational overhead.

Finally, the effectiveness of XAI for building trust in LMIC healthcare settings has not been studied yet.

4. A Conceptual AI-XR Framework for LMICs

The selection of generative AI, LLMs, MAS, and XAI as the core pillars of our framework was informed by a systematic analysis of barriers to XR rehabilitation in LMICs, drawn from the literature reviewed in Sections 2 and 3. Table 1 summarizes the key characteristics of the different technologies and their potential impact for rehabilitation, Table 2 maps each barrier to the AI technology best positioned to address it, based on the technology's key capabilities as documented in the literature. The framework is designed to be modular and interoperable, allowing incremental deployment as infrastructure and capacity grow.

Table 1. Key characteristic of four AI technologies to address specific rehabilitation challenges in LMICs.

AI Technology	Key Capability	Rehabilitation Challenge in LMICs	Potential Impact	Critical Limitations
Generative AI	Dynamic creation of 3D assets, environments, sounds	High cost and effort for culturally adapted content; limited diversity of exercises	Scalable, localized content; infinite variations of therapeutic tasks; reduced development burden	Real-time 3D generation on edge devices is challenging; risk of unsafe/hallucinated content
Large Language Models (LLMs)	Natural dialog, personalized coaching, explanation	Shortage of therapists; lack of personalized feedback; low health literacy	24/7 virtual coaching; multilingual support; enhanced motivation; support for community health workers	Hallucinations can cause harmful advice; computational demands; unproven for motion correction
Multiagent Systems (MAS)	Social agents, collaborative tasks, group dynamics	Social isolation; absence of group therapy; limited family involvement	Virtual peer support; collaborative rehabilitation; family engagement; scalable group sessions	Emergent behaviors unpredictable; psychological risks; limited clinical evidence
Explainable AI (XAI)	Transparent decision-making, visual rationale	Skepticism of AI; need for clinician oversight; regulatory requirements	Trustworthy automation; informed clinical supervision; bias detection; regulatory compliance	Explanations can be gamed; false reassurance; added complexity

Table 2. Barrier-to-technology mapping justifying the four AI pillars.

Barrier	AI Technology	Rationale
Lack of culturally adapted content	Generative AI	Automates creation of culturally specific 3D assets and environments [33]
Shortage of therapists for personalized coaching	LLMs	Enables conversational guidance and personalized feedback [38,39]
Social isolation and lack of group therapy	MAS	Provides virtual peers and collaborative tasks [25,45]
Mistrust of AI and need for clinical oversight	XAI	Makes decision-making transparent and auditable [26]
Limited language coverage	LLMs + Generative AI	LLMs for dialog, generative AI for voice synthesis in local languages [34,40]
Manual exercise progression	LLMs (analysis) + XAI (explanation)	Automated progression with transparent rationale
High development cost for new content	Generative AI	Reduces manual asset creation effort

4.1. Interaction Layer (LLMs + MAS)

The interaction layer manages the patient's real-time experience. It consists of two intertwined components: a conversational agent powered by LLMs and a multiagent system that populates the virtual world with social agents. The conversational agent acts as a virtual coach, using sensor data (e.g., hand tracking, joint angles) to provide feedback, answer questions, and offer encouragement. It can also explain the rationale behind exercises, draw analogies to daily activities, and adjust its tone based on the patient's emotional cues.

Given the risk of LLM hallucinations noted in Section 3.2, the conversational agent is constrained to a curated set of rehabilitation-specific response templates, with open-ended LLM generation used only for non-critical social dialog (e.g., encouragement, small talk). Clinical instructions and movement corrections are either template-based or require explicit clinician approval before deployment (in supervised mode).

The multiagent system generates virtual peers, coaches, and family members that interact with the patient in culturally appropriate ways. These agents can demonstrate exercises, engage in cooperative tasks (e.g., passing a ball), and provide social reinforcement. They can be configured to respect local norms—for instance, a virtual elder might speak more formally, while a peer agent uses colloquial language. The agents are also aware of the patient's performance and can adapt their behavior to provide appropriate challenges or support.

4.2. Decision Support Layer (XAI + Data Analytics)

The decision support layer processes the vast amount of data generated during therapy sessions to provide actionable insights for both patients and clinicians. Motion data (e.g., joint angles, movement smoothness, task completion times) are analyzed by AI models that track progress and detect compensatory movements. XAI modules then generate visual explanations that are presented to the user in an understandable format.

For patients, explanations might take the form of simple graphs showing improvement over time, accompanied by text like "Your shoulder range increased by 20%—you're ready for the next level." For clinicians, a dashboard provides detailed analytics, including progress charts, risk flags (e.g., compensation patterns that could lead to injury), and AI-generated recommendations for adjusting the therapy plan. All AI-generated recommendations should be explicitly labeled as "suggestions requiring clinical review" rather than automated decisions. The system also supports remote monitoring, allowing clinicians to review data and communicate with patients via the platform.

The decision support layer is built on a secure, privacy-preserving architecture, with data encrypted both in transit and at rest. Patient data is stored locally whenever possible, with only anonymized summaries sent to the cloud for aggregation and model improvement, in compliance with local data protection regulations [55].

4.3. Integration and Adaptability

The three layers are integrated through a shared data model and an event-driven architecture that ensures real-time responsiveness (Figure 1). The system is designed to be adaptable: new AI models can be swapped in as they become available, and local developers can extend the system with additional cultural assets or specialized modules. This flexibility is essential for long-term sustainability in dynamic African healthcare environments. Data flow diagram is presented in Figure 2.

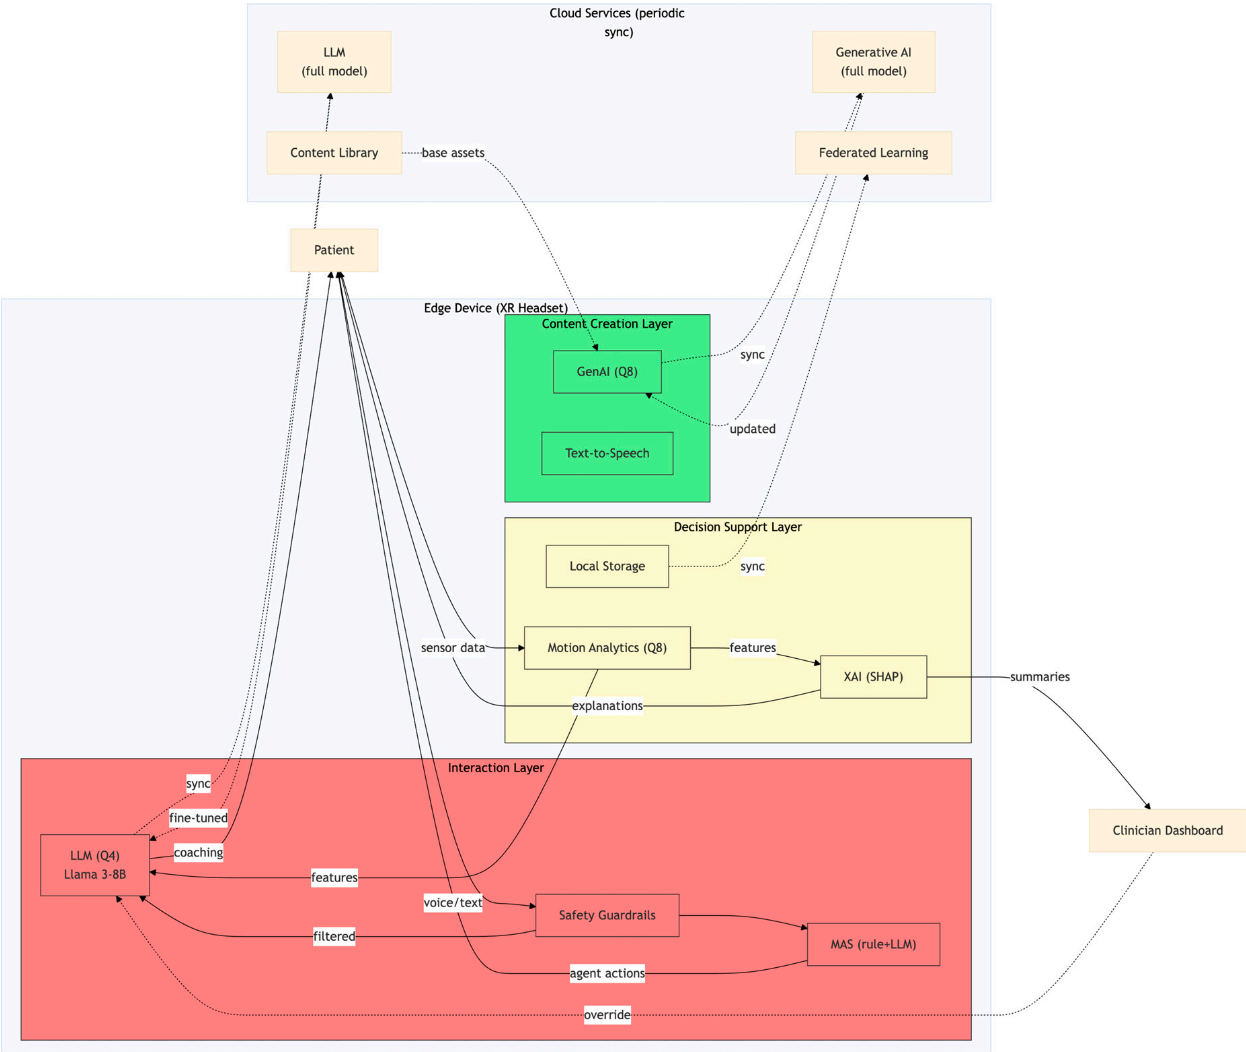

Figure 1. Detailed architecture of the proposed AI-enhanced XR rehabilitation system. The diagram shows the three layers (content creation, interaction, decision support), the separation between edge (XR headset/local device) and cloud, data flow pathways (solid = real-time, dashed = periodic sync), quantization points (Q4, Q8), and interactions among the four AI pillars (Generative AI, LLM, MAS, XAI). External entities (clinician dashboard, content library, patient) are also illustrated.

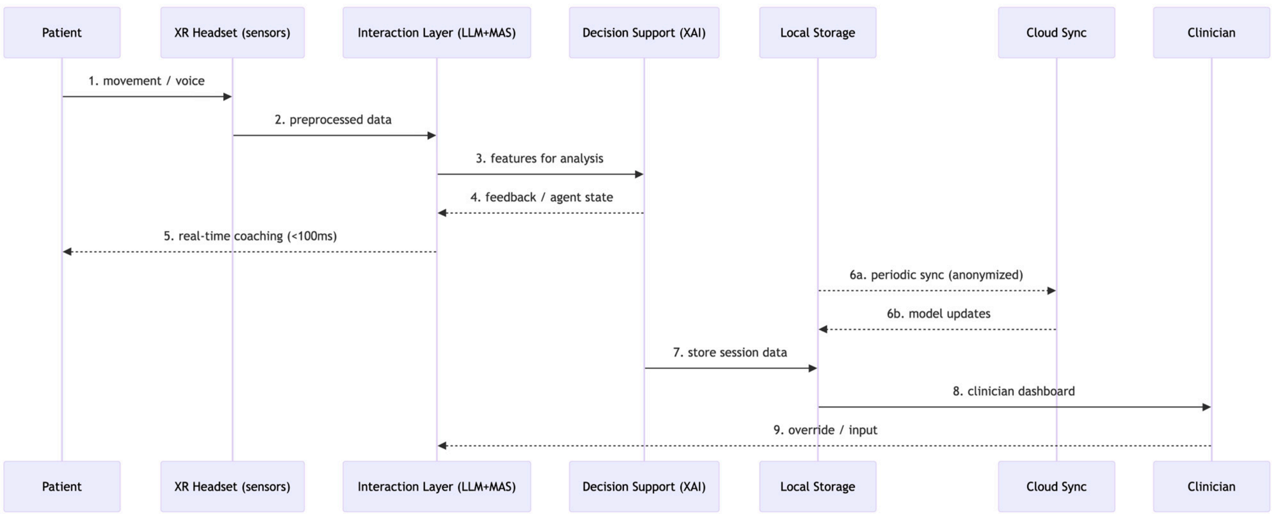

Figure 2. Sequence diagram illustrating end-to-end data flow and decision pathways in the AI-enhanced XR system. The diagram distinguishes between real-time (sub 100 ms) feedback loops, asynchronous cloud synchronization, and clinician oversight actions.

5. Implementation Considerations in African Contexts

Realizing the promise of AI-enhanced XR in Africa requires navigating a complex landscape of infrastructure, policy, and human factors. We highlight key considerations across technological, capacity, cultural, regulatory, and sustainability dimensions.

5.1. Technological Infrastructure

Internet connectivity remains uneven across Africa, but mobile penetration is high and growing, with a 4.6% annual growth in mobile subscriptions between 2019 and 2025 [56]. For XR applications, offline functionality is essential: the system must be able to run locally on headsets or tablets, with periodic cloud synchronization for updates and aggregated analytics. Generative AI and LLMs can be optimized to run on edge devices through techniques such as model quantization, pruning, and distillation, reducing dependence on continuous internet [57].

For deployment in rural settings, we propose three tiers of hardware configuration:

- Low-end: Standalone XR headset (e.g., Meta Quest 2) with quantized Llama 3-8B (4-bit, ~4 GB), template-based procedural content (no on-device generative AI), simple XAI (feature importance only);
- Mid-range: As above plus distilled diffusion model for 2D asset generation, more complex XAI (counterfactual explanations);
- High-end: Cloud-connected with full generative AI capabilities, real-time LLM, full MAS simulation (requires reliable 4G/5G).

Power stability is another challenge. In many rural areas, electricity supply is unreliable, with frequent load shedding [58]. The system should be energy-efficient and support solar charging or battery operation. Low-power XR headsets and the use of mobile phones as fallback devices for audio coaching can help mitigate power constraints.

5.2. Local Capacity Building

The co-creation model used in AdaptRehab VR should be extended to AI integration. Local rehabilitation professionals, software developers, and community health workers must be trained to configure, maintain, and interpret AI-enhanced systems. This includes understanding the capabilities and limitations of AI, how to use XAI outputs to inform clinical decisions, and how to troubleshoot common issues [18].

North–South training partnerships can play a vital role, but the goal must be to build self-sustaining local expertise [59–61]. This can be achieved through train-the-trainer programs, integration of AI and XR topics into university curricula in LMICs, and the creation of open educational resources [62]. Moreover, involving local developers in the co-creation of AI models ensures that the technology reflects local needs and values.

5.3. Cultural and Linguistic Adaptation

While generative AI can produce localized content, the underlying models must be trained on or fine-tuned with data that reflect African languages, cultural references, and social norms. This requires collaborative efforts to create open datasets—for example, multilingual rehabilitation dialogs, 3D models of everyday African objects, and culturally relevant exercise demonstrations. Such datasets are scarce; initiatives like Masakhane (for African languages) and local community archives can help fill this gap [40].

Moreover, cultural adaptation goes beyond language. It includes understanding local beliefs about health and disability, family dynamics, and the role of community in care [30]. Co-design sessions with end-users remain essential to ensure that the AI's behavior—such as the tone of the virtual coach or the social roles of agents—aligns with cultural expectations.

5.4. Ethical and Regulatory Frameworks

AI in healthcare raises concerns about data privacy, algorithmic bias, and autonomy. In LMICs, regulatory frameworks for AI-based medical devices are often nascent or nonexistent [52]. Ethical guidelines must be co-developed with local communities and health authorities, ensuring that data ownership remains with patients and that AI decisions do not reinforce existing disparities. Section 6 provides a comprehensive ethical analysis, including governance structures, accountability, liability, and equitable access. XAI becomes crucial here: explainability can help mitigate the "black box" problem [63] and facilitate regulatory review by making the system's logic transparent [38,64].

Informed consent processes should be adapted to local literacy levels and cultural contexts, explaining in plain language how AI is used and what data are collected. Special attention should be paid to the potential for AI to inadvertently exacerbate inequities if models are trained predominantly on data from one demographic group. Continuous monitoring for bias and regular model updates with diverse local data are essential practices [65].

5.5. Sustainability and Financing

The initial cost of XR hardware and AI development can be a barrier. Strategies to overcome this include leveraging open-source AI models, using affordable standalone headsets (e.g., Oculus Quest), and integrating with government health budgets or donor-funded programs. The return on investment—reduced reliance on expensive specialists, improved functional outcomes, lower long-term disability costs, and increased workforce productivity—should be systematically evaluated to build a business case [66].

Public–private partnerships may also play a role, with technology companies providing hardware or cloud services at reduced cost in exchange for research collaborations or social impact branding. Community-based health insurance schemes could cover AI-XR rehabilitation as a benefit, spreading costs across a larger population [67]. Long-term sustainability will require embedding the system into national rehabilitation strategies and securing ongoing funding for maintenance and updates.

https://doi.org/10.3390/app16104946

5.6. Technical Specifications and Benchmarking Roadmap

Full implementation details are beyond the scope of this conceptual paper. However, to guide future development, we propose the following specifications and benchmarking targets (Table 3), comparison with current system and the proposed solution is presented in Table 4.

Table 3. Proposed technical specifications (illustrative, not prescriptive).

Component	Model/Approach	Quantization	Memory	Latency Target
LLM (conversational)	Llama 3-8B or TinyLlama-1.1B	4-bit (INT4)	~4 GB or ~1 GB	<500 ms per response
Generative AI (3D)	Stable Diffusion XL + 3D uplift or procedural templates	8-bit	~2 GB	<10 s per asset (async)
XAI (feature importance)	SHAP with kernel approximation	n/a	~500 MB	<200 ms per explanation
MAS (agent simulation)	Rule-based + lightweight LLM for dialog	4-bit	~1 GB	<100 ms per agent update
Motion tracking	MediaPipe or custom CNN	8-bit	~200 MB	<50 ms per frame

Table 4. Current state versus the AI-enhanced vision across key dimensions of XR for rehabilitation.

Dimension	Current XR	AI-Enhanced XR (Proposed)
Content Creation	Manually designed 3D assets; limited cultural variation	Generative AI creates culturally adapted assets on demand (with limitations noted)
Language Support	Pre-recorded in one or few languages	LLMs support dozens of languages, with dialect adaptation (quality varies by language)
Feedback	Simple audio-visual cues and scores	Natural dialog, personalized coaching, movement corrections (safety constraints apply)
Social Interaction	Single-user, no virtual peers	Multiagent systems with virtual peers, coaches, family avatars (experimental)
Clinical Oversight	Manual progress tracking; limited explanation	XAI dashboards with visual rationales; automated suggestions (requiring clinician review)
Scalability	Each new language/culture requires separate development	Generative AI scales to new contexts with minimal manual effort (data availability permitting)
Autonomy	Dependent on therapist for customization	Adaptive AI personalizes therapy without constant human input (with clinician oversight)

Benchmarking roadmap (phased):

- Phase 1 (lab validation): Measure accuracy of LLM responses against gold-standard rehabilitation dialogs (target: >90% clinically appropriate); measure generative AI content safety (target: <1% unsafe/hallucinated assets); measure XAI explanation fidelity (target: >85% agreement with model internals).
- Phase 2 (field pilot): Deploy in 3–5 clinical sites across Africa; measure System Usability Scale (target: >70), task completion rates (target: >80%), user trust (adapted trust in automation scale).
- Phase 3 (comparative efficacy): Randomized controlled trial comparing AI-XR vs. conventional care; primary outcome: upper limb function (Fugl-Meyer); secondary: adherence, cost-effectiveness.

6. Ethical Considerations and Risks

While AI-enhanced XR holds great promise, it also introduces ethical risks that must be proactively addressed. These include potential over-reliance on technology, exacerbation of the digital divide, and unintended consequences of automated decision-making. We performed a PESTEL analysis to summarize the current situation and direction in Figure 3.

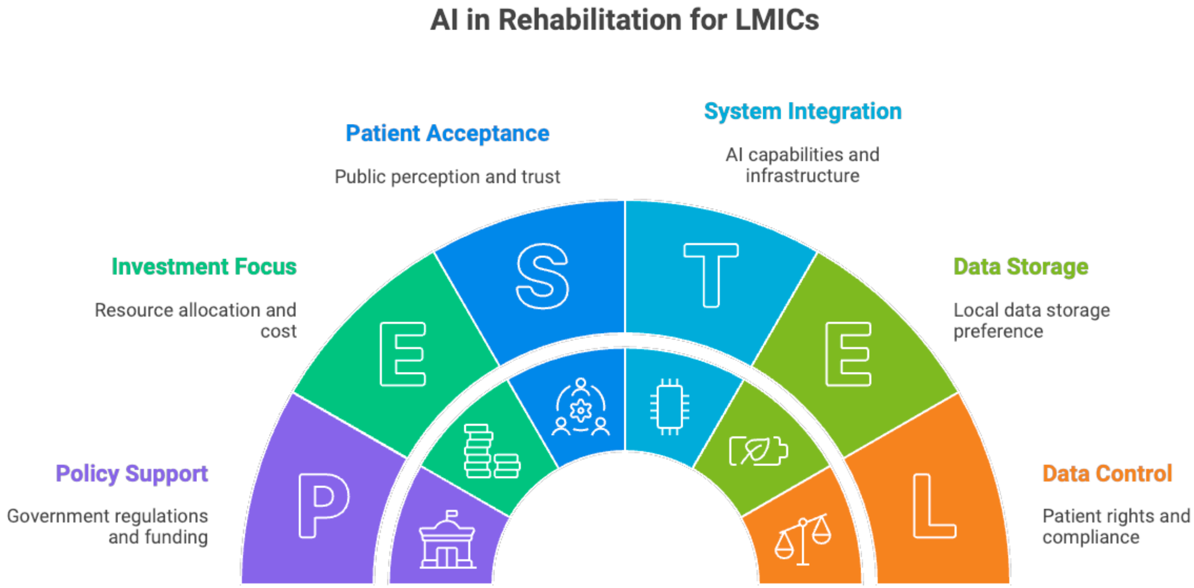

Figure 3. PESTEL (political, economic, social, technological, environmental, and legal) analysis of the use of AI in rehabilitation.

6.1. Autonomy and Human Oversight

There is a risk that AI systems could be seen as a replacement for human care, leading to reduced investment in rehabilitation professionals. However, in the LMIC context where the shortage is already severe, AI should be viewed as a complement that augments, not replaces, human capacity. We propose a "human-in-the-loop" model where: (i) all AI-generated treatment recommendations require clinician review before implementation; (ii) patients can opt out of AI coaching and receive simplified feedback; and (iii) an escalation pathway exists for cases where AI uncertainty is high. Clinical oversight must remain central: AI recommendations should be reviewed by a qualified professional whenever possible, and patients should have the option to speak with a human if they wish. XAI can support this by making the AI's reasoning transparent and allowing clinicians to override decisions when necessary.

6.2. Governance Structures

We propose a multi-stakeholder governance board for any AI-XR rehabilitation deployment, comprising:

- Clinical representatives (physiatrists, physical therapists) to ensure clinical safety;
- Patient and family advocates to represent user perspectives;
- Data protection officers to oversee privacy and data governance;
- AI ethics experts to monitor bias and fairness;
- Community health worker representatives to ensure frontline feasibility;
- Government health ministry officials to ensure alignment with national strategies.

The board's responsibilities would include: approving AI model updates before deployment; reviewing adverse event reports; conducting regular bias audits; managing patient data access requests; and publishing annual transparency reports.

6.3. Accountability and Liability

Determining accountability when an AI-XR system causes harm is complex. We propose a tiered liability framework:

- Hardware failures (e.g., XR headset malfunction): Manufacturer liability;
- AI model errors (e.g., incorrect exercise recommendation causing injury): Shared liability among developers and deploying institution, with clear contractual allocation;
- Clinical oversight failures (e.g., clinician ignored AI warning): Clinician/institution liability following standard medical malpractice principles;
- Patient misuse (e.g., ignoring safety warnings): Patient responsibility.

Regulatory bodies in LMICs should develop specific guidance for AI medical device liability. Pending such guidance, we recommend that deployments operate under research protocols with explicit informed consent and institutional indemnification.

6.4. Privacy and Data Security

Rehabilitation systems collect sensitive health data, including movement patterns that could reveal personal information. In contexts where data protection laws may be weak, robust technical safeguards are essential. This includes encryption, anonymization, and strict access controls. Data should be stored locally whenever possible, with only de-identified summaries shared for research or quality improvement. Patients should have control over their data and be informed about how it is used.

6.5. Algorithmic Bias and Fairness

AI models trained on data from high-income countries may not perform well on African populations due to differences in anthropometry, movement patterns, or language. Moreover, if data collection is skewed toward urban, educated populations, the models may be less accurate for rural or less literate users. It is crucial to involve diverse stakeholders in data collection and model validation, and to continuously monitor for performance disparities across demographic groups. XAI can help detect bias by revealing which features are driving decisions.

6.6. Informed Consent and Health Literacy

Obtaining meaningful informed consent for AI-enhanced rehabilitation requires careful adaptation. Patients may not understand what AI is or how it works, and traditional consent forms may be too complex. Interactive consent processes, using the XR environment itself to explain the system, could be developed. Visual explanations and simple analogies can help patients understand the benefits and risks, and they should be given the opportunity to ask questions.

6.7. Equitable Access and the Digital Divide

There is a risk that AI-XR rehabilitation could exacerbate existing health inequities if deployed only in urban, well-resourced settings. To promote equitable access, we recommend these points, which are summarized in Table 5.

- Tiered deployment strategy: Start with regional referral hospitals that have reliable infrastructure, then progressively decentralize to district hospitals, health centers, and community-based models;
- Shared device models: One XR headset per health facility, used for multiple patients on a rotating basis, rather than assuming individual ownership;
- Low-bandwidth fallbacks: Audio-only coaching via basic mobile phones for patients without XR access;

- Subsidized pricing: Differential pricing based on facility type (e.g., free for public health facilities, cost-recovery for private);
- Community health worker integration: Train CHWs to supervise AI-XR use, compensating them for this additional role to avoid burden shifting.

Table 5. Summary of ethical risks and mitigation strategies.

Risk	Mitigation Strategy
Over-reliance on AI, reduced human care	Human-in-the-loop; AI as complement, not replacement
Harmful AI recommendations	Output filtering; clinician review; safety guardrails
Data privacy breaches	Local storage; encryption; de-identification
Algorithmic bias	Diverse training data; regular bias audits; XAI for bias detection
Exacerbation of digital divide	Tiered deployment; shared devices; low-bandwidth fallbacks
Unclear liability	Tiered liability framework; research protocols; informed consent
LLM hallucinations	Constrained templates for clinical content; logging and audit
MAS emergent harmful behaviors	Simulation testing; human monitoring; kill switches
XAI false reassurance	Calibrated confidence displays; uncertainty quantification

7. Comparison with Existing Solutions and Value Proposition

To justify the proposed AI-XR framework, we compare it against existing rehabilitation delivery models in LMICs across six dimensions: cost-effectiveness, scalability, clinical evidence, cultural adaptability, personalization, and implementation complexity (Table 6).

Table 6. Comparative analysis of rehabilitation solutions.

Solution	Cost per Patient (Estimated)	Scalability	Clinical Evidence in LMICs	Cultural Adaptability	Personalization	Implementation Complexity
Conventional in-person therapy (therapist-led)	High ($100–500/session)	Low (therapist-limited)	Strong	High (human adaptation)	High	Moderate
Paper-based home exercise programs	Very low (<$1)	High	Moderate	Moderate (static)	Low	Low
Mobile health apps (2D)	Low ($5–20)	High	Emerging	Moderate (translatable)	Low-medium	Low
Non-AI XR	Medium ($300–800/device)	Medium (content limits)	Preliminary	High (co-created)	Medium (manual)	Medium
Telerehabilitation (remote clinician)	Medium ($50–150/session)	Low-medium (clinician-limited)	Moderate	Medium	High	Medium
Lower-cost 2D gamification (e.g., tablet-based)	Low ($50–200/device)	High	Emerging	Medium	Low-medium	Low
AI-XR (proposed)	High initially ($500–1000/device), but declining	Potentially high (AI reduces per-patient clinician time)	None yet (requires validation)	Potentially very high (generative adaptation)	Potentially very high (AI-driven)	High (AI integration, infrastructure)

The higher upfront cost and complexity of AI-XR are justified if it achieves three outcomes that lower-cost alternatives cannot: (i) engagement—maintaining adherence in unsupervised home settings through immersive, culturally adapted content; (ii) personalization—adapting to patient progress without requiring scarce clinician time; (iii) social support—addressing isolation through MAS, which is absent in other self-guided solutions. Whether

AI-XR delivers these benefits at acceptable cost is an empirical question requiring the trials proposed in Section 8.

8. Research Agenda and Future Directions

To translate the AI-enhanced XR vision into reality, a focused research agenda is needed. We propose the following priority areas, building on the foundations laid by previous LMIC-focused initiatives. A research agenda is presented in Table 7.

Table 7. Research agenda: priorities and timelines.

Timeframe	Priority	Description/Key Actions
Near-term (1–2 years)	1. Co-creation of datasets and pilot testing	• Create open datasets of African rehabilitation dialogs (minimum 10,000 utterances across 5 languages), 3D models of everyday African objects (minimum 500 objects), and culturally relevant exercise demonstrations. • Fine-tune open-source LLMs on rehabilitation dialogs and evaluate accuracy. • Pilot test a minimal viable AI-XR system (single AI pillar, e.g., LLM-only) with N = 20 patients to assess feasibility and safety.
	2. Technical optimization for edge deployment	• Benchmark quantized LLMs (3-bit, 4-bit, 8-bit) on representative XR hardware (e.g., Meta Quest 3) for inference latency and accuracy. • Develop lightweight procedural content generation as a fallback for generative AI. • Open-source the optimized models and evaluation code.
Medium-term (2–4 years)	3. Randomized controlled trials	• Conduct multi-site RCT comparing AI-XR vs. conventional • Measure functional outcomes, adherence, cost-effectiveness, and user trust. • Publish trial protocols and results with pre-specified analysis plans.
	4. Multilingual and multimodal expansion	• Expand language coverage to 10–15 African languages, including low-resource languages using transfer learning. • Integrate multimodal sensing (voice, movement, physiological) for more accurate patient state estimation. • Develop and validate XAI explanations for non-literate users (e.g., pictographic, audio).
Longer-term (3–5+ years)	5. Federated learning and continual adaptation	• Deploy federated learning to improve AI models across multiple sites without sharing raw patient data. • Develop online learning algorithms that personalize models to individual patients while maintaining safety. • Evaluate long-term outcomes (12+ months) and sustainability.
	6. Policy and scale-up	• Develop regulatory guidelines for AI medical devices in LMICs in partnership with health ministries. • Implement at national scale in 1–2 countries, measuring population-level impact. • Create open-source implementation toolkit to enable replication across Africa.

8.1. Co-Creation of AI Models with African Communities

Generative AI, LLMs, and MAS must be developed in partnership with local stakeholders to ensure cultural appropriateness and relevance. This involves creating open datasets of African languages, 3D objects, and rehabilitation scenarios, and fine-tuning models on these datasets. Research should explore participatory design methods that involve patients, clinicians, and community members in model development, validation, and iterative improvement.

8.2. Efficacy and Safety Trials

Randomized controlled trials are needed to compare AI-enhanced XR rehabilitation against conventional therapy and against non-AI XR. Outcomes should include functional measures (e.g., range of motion, activities of daily living), adherence, user satisfaction, and cost-effectiveness. Trials should be conducted in diverse settings (e.g., rural and urban, public and private) to assess generalizability. Special attention should be paid to safety, including the risk of falls or overexertion when patients use AI guidance without direct supervision.

8.3. Explainability and Trust Studies

Research should investigate how XAI influences trust and acceptance among African patients and clinicians. What formats of explanations (visual, textual, mixed) are most effective? Does XAI reduce technology anxiety and improve clinical decision-making? How can explanations be tailored to different literacy levels and cultural contexts? Longitudinal studies can assess whether trust evolves as users gain experience with the system.

8.4. Offline and Edge AI Optimization

Given infrastructure constraints, optimizing AI models to run on low-cost, offline-capable XR devices is critical. Research should focus on model compression, quantization, and efficient inference. Trade-offs between model size, accuracy, and energy consumption should be evaluated. Edge-cloud hybrid architectures that minimize data transfer while preserving functionality should be explored.

8.5. Implementation Science

Implementation studies should examine real-world deployment in diverse settings. What are the barriers and facilitators to adoption at the individual, organizational, and system levels? How can training programs for clinicians and community health workers be optimized? What financing models support long-term sustainability? Mixed-methods research can capture both quantitative outcomes and qualitative experiences.

8.6. Ethical Frameworks and Governance

Finally, research should contribute to the development of ethical guidelines and regulatory frameworks for AI-XR rehabilitation in LMICs. This includes best practices for data governance, informed consent, bias mitigation, and accountability. Engagement with policymakers and regulatory bodies is essential to ensure that frameworks are both protective and enabling.

9. Conclusions

The convergence of AI and XR holds transformative potential for rehabilitation in Africa, where the gap between need and available services is vast. Building on successful co-created XR systems like AdaptRehab VR, we have outlined how generative AI, large language models, multiagent systems, and explainable AI can together create a scalable, personalized, and culturally responsive rehabilitation ecosystem. Such a system can reduce dependence on scarce specialists, provide 24/7 support in local languages, foster social engagement, and maintain transparency to build trust.

However, this vision must be tempered with realism. The proposed AI technologies have significant limitations and failure modes documented in the literature; their effectiveness in African LMIC contexts is unproven; and infrastructure, capacity, and ethical challenges are substantial. Realizing the potential of AI-XR will require not only technical innovation but also deliberate efforts to co-design with African communities, invest in local capacity, adapt to infrastructure constraints, and develop ethical frameworks that ensure equitable benefits.

Realizing this vision requires deliberate efforts to co-design with African communities, invest in local capacity, adapt to infrastructure constraints, and develop ethical frameworks that ensure equitable benefits. The time is ripe to move from isolated XR applications to intelligent, AI-powered platforms that can democratize access to quality rehabilitation across the continent. Rigorous, context-sensitive research—beginning with the near-term priorities outlined above—is urgently needed to determine whether, how, and under what conditions AI-enhanced XR can deliver on its promise. By bridging the gap between technological innovation and local realities, AI-enhanced XR can help ensure that every stroke survivor, regardless of where they live, has a path to recovery.

Author Contributions: Conceptualization, C.D.K. and B.B.; Writing—original draft, C.D.K. and B.B.; Writing—review and editing, C.D.K. and B.B.; Visualization, B.B.; Supervision, B.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the NASCERE project (joint PhD program between Jimma University and Hasselt University) funded by VLIR-UOS. The funder has no role in the writing of the manuscript.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Johnston, S.C.; Mendis, S.; Mathers, C.D. Global Variation in Stroke Burden and Mortality: Estimates from Monitoring, Surveillance, and Modelling. *Lancet Neurol.* **2009**, *8*, 345–354. [CrossRef] [PubMed]
2. Sun, H.; Ma, B.; Jin, C.; Li, Z.; Song, X.; Bu, Y.; Liu, T.; Han, X.; Yang, X.; Yang, K.; et al. Global, Regional, and National Burdens of Stroke in Children and Adolescents From 1990 to 2019: A Population-Based Study. *Stroke* **2024**, *55*, 1543–1553. [CrossRef]
3. Felter, C.E.; Zalewski, K.; Jermann, R.; Palmer, P.L.; Baier, A.E.; Falvey, J.R. Rural Health: The Dirt Road Less Traveled. *Phys. Ther.* **2022**, *102*, pzac112. [CrossRef] [PubMed]
4. Cieza, A.; Causey, K.; Kamenov, K.; Hanson, S.W.; Chatterji, S.; Vos, T. Global Estimates of the Need for Rehabilitation Based on the Global Burden of Disease Study 2019: A Systematic Analysis for the Global Burden of Disease Study 2019. *Lancet* **2021**, *396*, 2006–2017. [CrossRef]
5. WHO Rehabilitation. 2030. Available online: https://www.who.int/initiatives/rehabilitation-2030 (accessed on 27 November 2024).
6. WHO. *Fact Sheets: Rehabilitation*; WHO: Geneva, Switzerland, 2020.
7. Olana, D.D.; Abessa, T.G.; Lamba, D.; Triccas, L.T.; Bonnechere, B. Effect of Virtual Reality-Based Upper Limb Training on Activity of Daily Living and Quality of Life among Stroke Survivors: A Systematic Review and Meta-Analysis. *J. Neuroeng. Rehabil.* **2025**, *22*, 92. [CrossRef]
8. Kenea, C.D.; Abessa, T.G.; Lamba, D.; Bonnechère, B. Immersive Virtual Reality in Stroke Rehabilitation: A Systematic Review and Meta-Analysis of Its Efficacy in Upper Limb Recovery. *J. Clin. Med.* **2025**, *14*, 1783. [CrossRef]
9. Bonnechère, B. *Serious Games in Physical Rehabilitation: From Theory to Practice*; Springer International Publishing: Berlin/Heidelberg, Germany, 2018, ISBN 978-3-319-66121-6.

10. Laver, K.E.; Lange, B.; George, S.; Deutsch, J.E.; Saposnik, G.; Chapman, M.; Crotty, M. Virtual Reality for Stroke Rehabilitation. *Cochrane Database Syst. Rev.* **2025**, *6*, CD008349. [CrossRef]
11. Kiper, P.; Godart, N.; Cavalier, M.; Berard, C.; Cieślik, B.; Federico, S.; Kiper, A.; Pellicciari, L.; Meroni, R. Effects of Immersive Virtual Reality on Upper-Extremity Stroke Rehabilitation: A Systematic Review with Meta-Analysis. *J. Clin. Med.* **2023**, *13*, 146. [CrossRef]
12. Chen, J.; Or, C.K.; Chen, T. Effectiveness of Using Virtual Reality-Supported Exercise Therapy for Upper Extremity Motor Rehabilitation in Patients With Stroke: Systematic Review and Meta-Analysis of Randomized Controlled Trials. *J. Med. Internet Res.* **2022**, *24*, e24111. [CrossRef]
13. Dias, P.; Silva, R.; Amorim, P.; Lains, J.; Roque, E.; Pereira, I.S.F.; Pereira, F.; Santos, B.S.; Potel, M. Using Virtual Reality to Increase Motivation in Poststroke Rehabilitation. *IEEE Comput. Graph. Appl.* **2019**, *39*, 64–70. [CrossRef]
14. Kenea, C.D.; Abessa, T.G.; Lamba, D.; Bonnechère, B. AdaptRehab VR: Development of an Immersive Virtual Reality System for Upper Limb Stroke Rehabilitation Designed for Low- and Middle-Income Countries Using a Participatory Co-Creation Approach. *Bioengineering* **2025**, *12*, 581. [CrossRef]
15. Kumar, S.; Pandey, S.K.; Kumar, A.; Kumar, R.; Kumar, N.; Kumar, S.; Kumar, A. Evaluating the Impact of Virtual Reality on Cognitive Recovery in Stroke Patients: A Comparative Single-Centered Study in the Indian Context. *Indian J. Phys. Med. Rehabil.* **2025**, *35*, 1–6. [CrossRef]
16. Lima, I.D.P.; Da Silva, M.B.B.; Fidelis, A.; Martins, B.; Leal, J.C.; Dos Santos Mendes, F.A. Usability and feasibility of immersive virtual games in the treatment of people with parkinson's disease. *Braz. J. Phys. Ther.* **2024**, *28*, 100666. [CrossRef]
17. Zheng, X.; Xing, L.; Lu, H.; Hao, S.; Liu, F. Traditional Rehabilitation Experiences, Unmet Needs, and Perspectives on Virtual Reality-Based Rehabilitation Among Patients with Stroke in China: Qualitative Thematic Analysis and Semistructured Interview Study. *J. Med. Internet Res.* **2026**, *28*, e84532. [CrossRef]
18. Dereje, D.; Lamba, D.; Abessa, T.G.; Kenea, C.; Ramari, C.; Osama, M.; Kossi, O.; Boma, P.M.; Panda, J.; Kushnir, A.; et al. Unlocking the Potential of Serious Games for Rehabilitation in Low and Middle-Income Countries: Addressing Potential and Current Limitations. *Front. Digit. Health* **2025**, *7*, 1505717. [CrossRef]
19. Rose, T.; Nam, C.S.; Chen, K.B. Immersion of Virtual Reality for Rehabilitation-Review. *Appl. Ergon.* **2018**, *69*, 153–161. [CrossRef]
20. Antonopoulos, P.; Fokides, E.; Koutromanos, G. Understanding Learning and Learning Experience in Immersive Virtual Reality. *Technol. Knowl. Learn.* **2025**, *30*, 1771–1800. [CrossRef]
21. Jesus, T.S.; Landry, M.D.; Dussault, G.; Fronteira, I. Human Resources for Health (and Rehabilitation): Six Rehab-Workforce Challenges for the Century. *Hum. Resour. Health* **2017**, *15*, 8. [CrossRef]
22. Kamalakannan, S.; Menon, B.; Solomon, J.M.; Musa, K.I. Editorial: Evidence on Low-Cost Technologies for Neurological Rehabilitation in Low and Middle-Income Countries. *Front. Neurol.* **2023**, *14*, 1323808. [CrossRef]
23. Sharma, W.; Lim, W.M.; Kumar, S.; Verma, A.; Kumra, R. Game on! A State-of-the-Art Overview of Doing Business with Gamification. *Technol. Forecast. Soc. Change* **2024**, *198*, 122988. [CrossRef]
24. Yang, X.; Chen, A.; PourNejatian, N.; Shin, H.C.; Smith, K.E.; Parisien, C.; Compas, C.; Martin, C.; Costa, A.B.; Flores, M.G.; et al. A Large Language Model for Electronic Health Records. *npj Digit. Med.* **2022**, *5*, 194. [CrossRef]
25. Dias, S.B.; Diniz, J.A.; Konstantinidis, E.; Savvidis, T.; Zilidou, V.; Bamidis, P.D.; Grammatikopoulou, A.; Dimitropoulos, K.; Grammalidis, N.; Jaeger, H.; et al. Assistive HCI-Serious Games Co-Design Insights: The Case Study of i-PROGNOSIS Personalized Game Suite for Parkinson's Disease. *Front. Psychol.* **2020**, *11*, 612835. [CrossRef]
26. Tjoa, E.; Guan, C. A Survey on Explainable Artificial Intelligence (XAI): Toward Medical XAI. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *32*, 4793–4813. [CrossRef]
27. Veras, M.; Labbé, D.R.; Furlano, J.; Zakus, D.; Rutherford, D.; Pendergast, B.; Kairy, D. A Framework for Equitable Virtual Rehabilitation in the Metaverse Era: Challenges and Opportunities. *Front. Rehabil. Sci.* **2023**, *4*, 1241020. [CrossRef]
28. Naderbagi, A.; Loblay, V.; Zahed, I.U.M.; Ekambareshwar, M.; Poulsen, A.; Song, Y.J.C.; Ospina-Pinillos, L.; Krausz, M.; Mamdouh Kamel, M.; Hickie, I.B.; et al. Cultural and Contextual Adaptation of Digital Health Interventions: Narrative Review. *J. Med. Internet Res.* **2024**, *26*, e55130. [CrossRef]
29. Fendt-Newlin, M.; Jagannathan, A.; Webber, M. Cultural Adaptation Framework of Social Interventions in Mental Health: Evidence-Based Case Studies from Low- and Middle-Income Countries. *Int. J. Soc. Psychiatry* **2020**, *66*, 41–48. [CrossRef]
30. Mueller, S.; Soriano, D.; Boscor, A.; Saville, N.; Arjyal, A.; Baral, S.; Fordham, M.; Hearn, G.; Le Masson, V.; Kayastha, R.; et al. MANTRA: Development and Localization of a Mobile Educational Health Game Targeting Low Literacy Players in Low and Middle Income Countries. *BMC Public Health* **2020**, *20*, 1171. [CrossRef]
31. Prvu Bettger, J.; Liu, C.; Gandhi, D.B.C.; Sylaja, P.N.; Jayaram, N.; Pandian, J.D. Emerging Areas of Stroke Rehabilitation Research in Low- and Middle-Income Countries: A Scoping Review. *Stroke* **2019**, *50*, 3307–3313. [CrossRef]
32. Li, X.; Elnagar, D.; Song, G.; Ghannam, R. Advancing Medical Education Using Virtual and Augmented Reality in Low- and Middle-Income Countries: A Systematic and Critical Review. *Virtual Worlds* **2024**, *3*, 384–403. [CrossRef]

33. Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; Ommer, B. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 18–24 June 2022*; IEEE: Piscataway, NJ, USA, 2022; pp. 10674–10685.
34. Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; Stoyanov, V. Unsupervised Cross-Lingual Representation Learning at Scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2020; pp. 8440–8451.
35. Mekbib, D.B.; Debeli, D.K.; Zhang, L.; Fang, S.; Shao, Y.; Yang, W.; Han, J.; Jiang, H.; Zhu, J.; Zhao, Z.; et al. A Novel Fully Immersive Virtual Reality Environment for Upper Extremity Rehabilitation in Patients with Stroke. *Ann. N. Y. Acad. Sci.* **2021**, *1493*, 75–89. [CrossRef]
36. Ouedraogo, I.; Some, B.M.J.; Oyibo, K.; Benedikter, R.; Diallo, G. Using Serious Mobile Games to Improve Health Literacy in Rural Sub-Saharan Africa: A Literature Review. *Front. Public Health* **2022**, *10*, 768252. [CrossRef]
37. Wang, Y.; Guizani, M.; Hossain, M.S. Artificial Intelligence for Virtual Reality: State of the Art, Challenges, and Future Perspectives. *ACM Trans. Multimed. Comput. Commun. Appl.* **2025**, *21*, 1–29. [CrossRef]
38. Qiang, S.; Zhang, H.; Liao, Y.; Zhang, Y.; Gu, Y.; Wang, Y.; Xu, Z.; Shi, H.; Han, N.; Yu, H. Application of Large Language Models in Stroke Rehabilitation Health Education: 2-Phase Study. *J. Med. Internet Res.* **2025**, *27*, e73226. [CrossRef]
39. Lee, P.; Bubeck, S.; Petro, J. Benefits, Limits, and Risks of GPT-4 as an AI Chatbot for Medicine. *N. Engl. J. Med.* **2023**, *388*, 1233–1239. [CrossRef]
40. Adelani, D.I.; Abbott, J.; Neubig, G.; D'souza, D.; Kreutzer, J.; Lignos, C.; Palen-Michel, C.; Buzaaba, H.; Rijhwani, S.; Ruder, S.; et al. MasakhaNER: Named Entity Recognition for African Languages. *Trans. Assoc. Comput. Linguist.* **2021**, *9*, 1116–1131. [CrossRef]
41. Bonnechère, B.; Kossi, O.; Mapinduzi, J.; Panda, J.; Rintala, A.; Guidetti, S.; Spooren, A.; Feys, P. Mobile Health Solutions: An Opportunity for Rehabilitation in Low- and Middle Income Countries? *Front. Public Health* **2022**, *10*, 1072322. [CrossRef]
42. Zhong, X.; Li, S.; Chen, Z.; Ge, L.; Yu, D.; Wang, S.; You, L.; Shang, H. Considerations for Patient Privacy of Large Language Models in Health Care: Scoping Review. *J. Med. Internet Res.* **2025**, *27*, e76571. [CrossRef]
43. Chiaramonte, R.; D'Amico, S.; Caramma, S.; Grasso, G.; Pirrone, S.; Ronsisvalle, M.G.; Bonfiglio, M. The Effectiveness of Goal-Oriented Dual Task Proprioceptive Training in Subacute Stroke: A Retrospective Observational Study. *Ann. Rehabil. Med.* **2024**, *48*, 31–41. [CrossRef]
44. Goumopoulos, C.; Chartomatsidis, M.; Koumanakos, G. Participatory Design of Fall Prevention Exergames Using Multiple Enabling Technologies. In *Proceedings of the 8th International Conference on Information and Communication Technologies for Ageing Well and e-Health, Online, 23–25 April 2022*; SCITEPRESS-Science and Technology Publications: Setubal, Portugal, 2022; pp. 70–80.
45. Afyouni, I.; Rehman, F.U.; Qamar, A.M.; Ghani, S.; Hussain, S.O.; Sadiq, B.; Rahman, M.A.; Murad, A.; Basalamah, S. A Therapy-Driven Gamification Framework for Hand Rehabilitation. *User Model. User-Adapt. Interact.* **2017**, *27*, 215–265. [CrossRef]
46. Bunt, L.; Greeff, J.; Taylor, E. Enhancing Serious Game Design: Expert-Reviewed, Stakeholder-Centered Framework. *JMIR Serious Games* **2024**, *12*, e48099. [CrossRef]
47. Vourvopoulos, A.; Faria, A.L.; Ponnam, K.; Bermudez i Badia, S. RehabCity: Design and Validation of a Cognitive Assessment and Rehabilitation Tool through Gamified Simulations of Activities of Daily Living. In *Proceedings of the 11th Conference on Advances in Computer Entertainment Technology*; Association for Computing Machinery: New York, NY, USA, 2014; pp. 1–8.
48. Rozevink, S.G.; van der Sluis, C.K.; Garzo, A.; Keller, T.; Hijmans, J.M. HoMEcare aRm rehabiLItatioN (MERLIN): Telerehabilitation Using an Unactuated Device Based on Serious Games Improves the Upper Limb Function in Chronic Stroke. *J. Neuroeng. Rehabil.* **2021**, *18*, 48. [CrossRef]
49. Russell, S.J.; Norvig, P. *Artificial Intelligence: A Modern Approach*, 4th ed.; Prentice Hall Series in Artificial Intelligence; Global Edition; Pearson: Boston, MA, USA, 2022; ISBN 978-1-292-40113-3.
50. Rahimi, B.; Nadri, H.; Lotfnezhad Afshar, H.; Timpka, T. A Systematic Review of the Technology Acceptance Model in Health Informatics. *Appl. Clin. Inform.* **2018**, *9*, 604–634. [CrossRef]
51. Kushnir, A.; Kachmar, O.; Bonnechère, B. STASISM: A Versatile Serious Gaming Multi-Sensor Platform for Personalized Telerehabilitation and Telemonitoring. *Sensors* **2024**, *24*, 351. [CrossRef]
52. Iyamu, I.; Gómez-Ramírez, O.; Xu, A.X.; Chang, H.-J.; Watt, S.; Mckee, G.; Gilbert, M. Challenges in the Development of Digital Public Health Interventions and Mapped Solutions: Findings from a Scoping Review. *Digit Health* **2022**, *8*, 20552076221102255. [CrossRef] [PubMed]
53. Rudin, C. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215. [CrossRef]
54. Bansal, G.; Nushi, B.; Kamar, E.; Lasecki, W.S.; Weld, D.S.; Horvitz, E. Beyond Accuracy: The Role of Mental Models in Human-AI Team Performance. *Proc. AAAI Conf. Hum. Comput. Crowdsourc.* **2019**, *7*, 2–11. [CrossRef]
55. WHO. *WHO Guideline: Recommendations on Digital Intervention for Health System Strenghtening*; WHO: Geneva, Switzerland, 2019.

56. GSMA. *The Mobile Economy Sub-Saharan Africa 2024*; GSMA: Nairobi, Kenya, 2024.
57. Deng, B.L.; Li, G.; Han, S.; Shi, L.; Xie, Y. Model Compression and Hardware Acceleration for Neural Networks: A Comprehensive Survey. *Proc. IEEE* **2020**, *108*, 485–532. [CrossRef]
58. Bakht, M.P.; Salam, Z.; Bhatti, A.R.; Ullah Sheikh, U.; Khan, N.; Anjum, W. Techno-Economic Modelling of Hybrid Energy System to Overcome the Load Shedding Problem: A Case Study of Pakistan. *PLoS ONE* **2022**, *17*, e0266660. [CrossRef]
59. Maponga, C.C.; Mhazo, A.T.; Morse, G.D. A Framework for Sustainable Capacity-Building for Collaborative North-South Translational Health Research and Training in a Resource-Constrained Setting. *Health Res. Policy Syst.* **2023**, *21*, 24. [CrossRef]
60. Clark, T.; Verastegui, M.; Tinajeros, F.; Calderon, M.; Colanzi, R.; Gilman, R.H. Building Public Health Capacity through a Sustainable South-South-North Training Program. *Am. J. Trop. Med. Hyg.* **2020**, *103*, 2581–2583. [CrossRef] [PubMed]
61. Bonnechère, B. Building Equitable Health Partnerships: Addressing Racial Disparities in Global Health. *Front. Public Health* **2025**, *13*, 1604892. [CrossRef]
62. O'Brien, P.; Kajja, I.; Potter, J.M.; O'Hara, N.N.; Kironde, E.; Petrisor, B. Role of North-South Partnership in Trauma Management: Uganda Sustainable Trauma Orthopaedic Program. *J. Orthop. Trauma* **2018**, *32*, S21–S24. [CrossRef] [PubMed]
63. Bonnechère, B. Unlocking the Black Box? A Comprehensive Exploration of Large Language Models in Rehabilitation. *Am. J. Phys. Med. Rehabil.* **2024**, *103*, 532–537. [CrossRef] [PubMed]
64. Neo, J.R.E.; Ser, J.S.; Tay, S.S. Use of Large Language Model-Based Chatbots in Managing the Rehabilitation Concerns and Education Needs of Outpatient Stroke Survivors and Caregivers. *Front. Digit. Health* **2024**, *6*, 1395501. [CrossRef]
65. Obermeyer, Z.; Powers, B.; Vogeli, C.; Mullainathan, S. Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science* **2019**, *366*, 447–453. [CrossRef]
66. Ramos, Q.M.R.; Neill, R.; Bachani, A.M.; Trujillo, A.J. The Return of Investing Healthcare Resources in Rehabilitation: Should Governments Push the Pedal or Slowdown? *Res. Sq.* 2024, *preprint*.
67. Labrique, A.B.; Wadhwani, C.; Williams, K.A.; Lamptey, P.; Hesp, C.; Luk, R.; Aerts, A. Best Practices in Scaling Digital Health in Low and Middle Income Countries. *Glob. Health* **2018**, *14*, 103. [CrossRef] [PubMed]

https://doi.org/10.3390/app16104946