

**FROM ISOLATED SYSTEMS TO DISTRIBUTED INTELLIGENCE: AN EVOLUTIONARY AND CRITICAL REVIEW OF PERSONALIZED REHABILITATION**

Hasna K H	Research Scholar, PG & Research Department of Computer Science, Nilgiri College of Arts and Science(Autonomous), Thaloor, The Nilgiris, Tamil Nadu
Dr. T. Preethi	Assistant Professor, PG & Research Department of Computer Science, Nilgiri College of Arts and Science(Autonomous), Thaloor, The Nilgiris, Tamil Nadu
Mr. Muhammad Savad N	Research Scholar, PG & Research Department of Computer Science, Nilgiri College of Arts and Science(Autonomous), Thaloor, The Nilgiris, Tamil Nadu

ABSTRACT Rehabilitation technology has evolved through a sequence of increasingly capable relationships between the patient, therapist, machine, and data. Mechanical assistance first improved repetition and consistency; sensing then made movement measurable; artificial intelligence enabled interpretation and prediction; adaptive control introduced responsiveness; and connected systems expanded rehabilitation beyond a single device. This review traces that progression and examines the emerging transition toward distributed intelligence, in which multiple sensing and robotic systems may contribute complementary information while learning collaboratively. The scope is deliberately narrowed only after the broader evolution is established. The later sections examine pediatric rehabilitation, where growth, developmental variability, engagement, individualized goals, caregiver involvement, and physical safety create requirements that cannot be assumed from adult rehabilitation evidence. The review synthesizes literature on robotic gait and upper-limb rehabilitation, wearable and sensor-assisted systems, artificial intelligence, computer vision, adaptive and reinforcement learning, socially assistive robotics, multimodal and multi-robot systems, edge intelligence, federated learning, explainable AI, privacy, security, and clinical translation. Across these areas, the literature demonstrates substantial progress in individual technological components but less mature integration across systems. The strongest evidence concerns repetitive robotic assistance and quantitative assessment, whereas external validation, long-term personalization, interoperable multi-robot intelligence, pediatric federated learning, clinically useful explainability, and longitudinal real-world benefit remain comparatively limited. The review therefore identifies a transition point rather than presenting a completed architecture: future rehabilitation research may need to move from isolated intelligent devices toward distributed, privacy-conscious, child-centered environments operating under explicit clinical oversight. The final sections derive research gaps and a staged roadmap for evaluating this direction without equating technological novelty with clinical effectiveness.

KEYWORDS : Rehabilitation Systems, Intelligent Rehabilitation, Personalized Rehabilitation, Artificial Intelligence, Rehabilitation Robotics, Pediatric Rehabilitation, Edge Intelligence, Federated Learning, Explainable AI, Multi-Robot Systems, Clinical Decision Support

I. INTRODUCTION

Rehabilitation has traditionally relied on therapist-guided, repetitive, task-oriented practice. This model remains essential, but its effectiveness is influenced by the amount of practice that can be delivered, the consistency of repetition, the ability to quantify progress, and the time available for individualized adjustment. These constraints have encouraged a gradual shift from purely manual treatment toward technological assistance.

The initial response was mechanical. Devices could reproduce trajectories, provide controlled forces, or support a limb during repetitive exercise. This improved consistency, but the device generally had limited awareness of the person's changing state. The next response was therefore to add sensing. Once rehabilitation systems could measure movement, force, pressure, position, physiological signals, or interaction, they could begin to quantify what was happening during therapy.

Measurement created a new problem: rehabilitation data are complex, continuous, and highly variable. Artificial intelligence subsequently became useful for classifying movement, predicting outcomes, recognizing intention, and identifying patterns that are difficult to capture with simple rules. Adaptive control then extended this capability by allowing assistance to change according to observed performance.

The evolution did not stop at a single intelligent device. Rehabilitation increasingly involves combinations of robots, wearables, cameras, edge devices, and clinical information systems. This creates a further challenge: how can multiple sources of intelligence work together without producing fragmented information, excessive latency, or unacceptable privacy and security risks?

Only at this point does the pediatric context become central. Children introduce additional requirements because growth, developmental variation, motivation, communication, caregiver participation, and safety alter both the meaning of rehabilitation data and the design of physical interaction. Pediatric rehabilitation therefore provides a

demanding test case for the broader evolution toward intelligent and distributed rehabilitation.

This review follows that evolution deliberately. Rather than announcing a specific multi-robot or federated architecture at the outset, it examines why each technological transition occurred, what limitation remained, and how the next generation attempted to address it. The progression of the literature thus informs the research gaps identified near the end of the review.

II. Review Scope and Analytical Frame

This manuscript is a structured evolutionary and critical review based on the curated literature set supplied for the study. The set contains 20 papers spanning rehabilitation robotics, pediatric robotic design, gait rehabilitation, artificial intelligence and machine learning, adaptive control, social robotics, federated learning, explainable AI, edge intelligence, multi-agent robotics, robot learning, clinical decision support, and digital twins.

The review is organized around a causal sequence rather than a device catalog. Each stage asks four questions: What capability did the existing system provide? What limitation remained? Why did the next technological development become necessary? And what new problem did that development introduce?

The evidence is considered across five dimensions: technological capability, pediatric suitability, clinical evidence, data and AI requirements, and translational readiness. This approach prevents technical performance from being treated as synonymous with clinical effectiveness.

The supplied literature set does not include a complete reproducible database-search history, duplicate-removal log, and study-level screening record for all included papers. Therefore, the present manuscript is explicitly positioned as an evolutionary and critical narrative review rather than a formal PRISMA systematic review. PRISMA 2020 is recognized as an appropriate reporting framework for a future systematic extension, but no screening counts are fabricated here.

Table I. Causal Evolution of Rehabilitation Technology

Stage	Capability gained	Limitation revealed	Next transition
Mechanical assistance	Repeatable movement and controlled support	Limited awareness and adaptation	Sensing
Sensor-assisted systems	Objective movement and force measurement	Large, complex data	AI interpretation
AI-assisted systems	Classification, prediction, recognition	Limited generalization and personalization	Adaptive learning
Adaptive systems	Performance-responsive assistance	Safety and clinical validation	Multimodal integration
Connected systems	Multiple complementary data sources	Interoperability and latency	Edge/distributed intelligence
Distributed learning	Cross-site collaborative learning	Heterogeneity, security, explainability	Trustworthy clinical intelligence
Clinical intelligence	Decision support under human oversight	Need for prospective validation	Translation

III. From Conventional Rehabilitation to Mechanical Assistance

Why was mechanical assistance needed?

Conventional rehabilitation can provide individualized human judgment, but high-repetition treatment places practical demands on therapists and may make continuous quantitative assessment difficult. Mechanical assistance emerged to make selected therapeutic movements more repeatable and controllable.

Mechanically assisted rehabilitation systems initially emphasized controlled trajectories, physical support, and repeated practice. Their principal contribution was not intelligence but consistency. By reducing variation in how a movement works, these systems created a more controlled environment for rehabilitation research and treatment, these systems created a more controlled environment for rehabilitation research and treatment.

The limitation was equally important. A device could reproduce a movement without necessarily understanding whether the patient was improving, compensating, becoming fatigued, or requiring a different level of assistance. This limitation created the motivation for rehabilitation robotics and sensor integration.

Pediatric design makes the mechanical problem more demanding. Gonzalez et al. emphasized pediatric design features such as safety, weight, operability, and adaptability, demonstrating that a pediatric device cannot simply be treated as a smaller version of an adult system [2].

The next question, therefore, became: If a machine is assisting movement, how can it know what the person is actually doing?

IV. From Mechanical Assistance to Robotic and Sensor-driven Systems

Robotic rehabilitation expanded mechanical assistance by combining actuation with sensing and control. Gait trainers, exoskeletons, treadmill systems, upper-limb robots, robotic walkers, and other assistive platforms can deliver structured movement while recording kinematic and kinetic information.

Gait and lower-limb rehabilitation is among the most established applications. Robotic gait systems can support repetitive stepping while measuring movement and interaction. Upper-limb systems similarly support reaching, grasping, coordination, and task-oriented practice while generating measures such as trajectory deviation, range of motion, force, smoothness, and task completion.

Wearable and assistive robotics extend this capability beyond a fixed workstation. Robotic walkers, exoskeletons, powered mobility devices, prosthetic systems, and robotic arms can support independence and activities of daily living. Socially assistive robots extend the role of robotics further by addressing motivation, attention, communication, and behavioral participation.

The key transition is from a machine that merely delivers assistance to a system that observes the interaction. Force, position, inertial,

pressure, electromyographic, and vision sensors allow the system to estimate movement quality, effort, symmetry, and interaction state.

However, sensing introduces a new limitation: data volume and complexity. Once rehabilitation becomes measurable, the challenge shifts from collecting data to interpreting it. This creates the logical entry point for artificial intelligence.

The field therefore moved from asking “How can the robot assist?” to asking “What does the collected information mean?”

V. From Measurement to Artificial Intelligence

Machine learning and deep learning have been applied to movement classification, gait analysis, trajectory prediction, intention recognition, outcome assessment, and adaptive control. AI is particularly useful when rehabilitation signals are multidimensional and when clinically meaningful patterns are difficult to express as fixed rules.

Computer vision enables markerless estimation of posture, joint configuration, facial expression, and movement quality. This can reduce sensor burden, but performance remains sensitive to occlusion, lighting, camera position, clothing, and changes between laboratory and real-world environments.

Reinforcement learning and policy-learning approaches introduce another step by allowing a system to select actions based on observed performance. Reviews of intelligent control in rehabilitation describe the movement toward adaptive strategies rather than fixed trajectories [7]. Yet physical rehabilitation is not a suitable environment for unconstrained trial-and-error; safety constraints and bounded exploration are essential.

Natural language processing and affective computing broaden intelligence beyond movement. Instruction, feedback, emotional state recognition, and social interaction can influence engagement. Socially assistive robotics is therefore part of the same evolutionary trajectory: the system is no longer only controlling movement but interpreting aspects of human behavior.

The pediatric AI literature reinforces the potential of this transition while also exposing its limitations. Kaelin et al. reported extensive AI use in rehabilitation targeting children and youth with disabilities but found limited individualized participation-oriented personalization [5]. More recent pediatric AI outcome-assessment work similarly highlights methodological heterogeneity, small samples, and limited external validation [6].

AI therefore solves the interpretation problem but creates a new question: if a system can understand performance, can it use that information to provide the right level of assistance for an individual?

VI. From Artificial Intelligence to Adaptive and Personalized Rehabilitation

Personalization is frequently described as the defining promise of AI-enabled rehabilitation. In practice, however, personalization can refer to very different levels of adaptation. A system may change task difficulty, assistance level, feedback frequency, or exercise selection without actually connecting those changes to a child's long-term functional goals.

Children with the same diagnosis can differ substantially in age, developmental stage, motor capacity, motivation, fatigue, communication, and therapeutic objectives. Consequently, population-level protocols may be useful starting points but cannot fully represent individual rehabilitation trajectories.

The distinction between adaptation and personalization is important. Adaptation is a system response to observed performance. Personalization is a clinically meaningful process in which the system's behavior is linked to the individual's characteristics, goals, progression, and context. The pediatric AI scoping evidence suggests that much existing work remains closer to group-level tailoring than truly individualized goal-setting [5].

Longitudinal personalization also requires stable outcome measurement. If a system continuously changes the therapy protocol, raw scores cannot be compared without considering task difficulty and assistance level. Future systems therefore need longitudinal patient representations that preserve context rather than isolated session scores.

This approach creates another limitation. Even a highly personalized single robot can observe only a subset of rehabilitation. Gait, upper-limb coordination, activity outside therapy, engagement, and physiological responses may require different sensors or devices.

Personalization therefore leads naturally to the next question: can multiple specialized systems contribute to a shared understanding of rehabilitation?

VII. From Individual Intelligence to Multimodal and Multi-Robot Systems

No single rehabilitation device observes the entire patient. A gait robot captures lower-limb function; an upper-limb robot captures reaching and coordination; wearables provide continuous movement information; cameras provide visual observations; and socially assistive robots can capture interaction and engagement. Each source is incomplete but potentially complementary.

Multimodal intelligence attempts to combine these sources. Its value lies in the possibility that one signal can provide context for another. Its difficulty lies in synchronization, missing data, different sampling rates, incompatible coordinate systems, inconsistent data formats, and differences in measurement reliability.

Multi-robot rehabilitation extends the idea from multimodal sensing to coordinated physical systems. Specialized robots can be treated as complementary agents rather than independent devices. However, coordination requires more than communication. The systems need shared representations, interoperable data, compatible models, and mechanisms for resolving conflicting observations.

Jiang et al. provide an important proof of concept for collaborative learning across rehabilitation robots using federated joint learning in stroke rehabilitation [6]. This demonstrates that rehabilitation robots can participate in collaborative model learning, but it should not be interpreted as pediatric clinical evidence.

The broader literature therefore suggests that multi-robot intelligence is technically plausible but clinically fragmented. The field has multiple intelligent components, yet these components rarely form a common longitudinal representation of the patient.

Once multiple systems are connected, another question becomes unavoidable: where should their computation occur, and how should their data be shared?

VIII. From Connected Systems to Edge Intelligence

Rehabilitation robots and wearables generate continuous movement, force, physiological, and interaction data. Sending every signal to a remote server may introduce latency, bandwidth requirements, and additional exposure of sensitive information.

Edge intelligence moves selected computation closer to the patient. Local preprocessing, feature extraction, anomaly detection, movement estimation, and time-sensitive inference can be performed on or near the rehabilitation device. This is particularly relevant for safety monitoring, where a robot should not depend entirely on a remote server before responding to an unsafe event.

Edge deployment nevertheless introduces constraints. Wearable and embedded platforms may have limited memory, processing power, battery capacity, and thermal headroom. Deep models may therefore need compression, efficient architectures, or task-specific deployment strategies.

Edge computing does not solve the institutional data-sharing problem by itself. It addresses where computation occurs, whereas distributed learning addresses how multiple sites can collaborate. These concepts therefore become complementary rather than competing.

The next problem is institutional: if useful rehabilitation intelligence requires data from multiple sites, how can learning occur without requiring all sensitive data to be centralized?

IX. From Edge Intelligence to Distributed and Federated Learning

Federated learning is attractive because rehabilitation data are distributed across hospitals, clinics, laboratories, and home environments. In a federated setting, local systems can train models and share model information rather than transferring raw patient data to a central repository.

The privacy benefit should not be overstated. Model updates can remain vulnerable to inference, poisoning, and other attacks. Secure aggregation, authentication, differential privacy, robust aggregation, and careful governance may be required.

Pediatric rehabilitation introduces an additional challenge: heterogeneous data. Children differ by age, diagnosis, developmental stage, functional severity, therapy protocol, sensor configuration, and institution. A model that performs well on average may therefore perform poorly for a particular subgroup or child.

Healthcare federated-learning literature identifies heterogeneity as a central challenge, while federated rehabilitation-robot learning demonstrates the technical feasibility of collaborative model training [6], [17], [18]. The research opportunity is to determine whether privacy-preserving collaboration can improve generalization without erasing clinically important local differences.

Federated learning should consequently be treated as an enabling mechanism rather than the final clinical goal. The relevant question is whether distributed collaboration produces better and safer rehabilitation intelligence than isolated local learning.

Distributed learning solves part of the data problem, but a new clinical question appears: can professionals understand and appropriately challenge the decisions produced by distributed AI?

X. From Intelligent Prediction to Trustworthy Clinical Intelligence

Explainable AI becomes important as AI outputs move from measurement toward recommendation. A therapist may need to know why a system identified a movement as abnormal, why assistance was increased, or why a child's performance was judged to have changed.

Medical XAI research emphasizes that explanation quality includes more than mathematical feature attribution. An explanation should be faithful to the model, useful to the intended user, and compatible with clinical workflow [19]. Engineers may prefer detailed feature maps or confidence distributions, whereas therapists may need a concise explanation linked to observable rehabilitation measures.

Privacy and explainability also need to be considered together in distributed learning. A model may be trained collaboratively without centralizing raw data while still producing outputs that are difficult to interpret. Trustworthy distributed intelligence therefore requires both data governance and human-understandable decision support.

Safety is particularly important when AI interacts physically with a patient. A classification error may affect documentation, but an incorrect robotic action can create physical risk. Mechanical limits, monitoring, emergency stops, therapist overrides, and explicit fallback behavior should therefore be treated as part of the intelligent system rather than as peripheral engineering details.

Pediatric data governance adds consent; assent, where applicable; data minimization; access control; retention; secondary use; and privacy considerations. Fairness is also important because datasets may underrepresent certain ages, diagnoses, functional abilities, socioeconomic groups, or geographic populations.

At this point, the broad technological evolution converges on a more demanding application context. The next section asks what changes when these intelligent systems are intended specifically for children.

XI. The Pediatric Context: Why the Requirements Become More Demanding

Pediatric rehabilitation is not simply adult rehabilitation performed with smaller equipment. Children grow, develop, learn, communicate, and engage with therapy differently. Their physical dimensions change over time, and their functional goals may extend from impairment reduction to school participation, play, mobility, communication, and independence.

A. Growth and Anthropometric Variability

Growth changes body dimensions, alignment, range of motion, and mechanical interaction with robotic devices. Systems must therefore accommodate changing physical characteristics rather than relying on fixed assumptions.

B. Developmental and Clinical Heterogeneity

The same diagnosis may produce substantially different movement

patterns and functional needs. Pediatric-specific datasets and external validation are needed for accurate classification of developmental variation vs misclassification of impairment by AI models.

C. Engagement and Motivation

Children's response to repetitive therapy can depend on attention, motivation, interaction, and perceived relevance. Socially assistive robots are relevant here, although the evidence remains heterogeneous. The pediatric social-robot literature illustrates the need for stronger evaluation of sustained engagement and clinical benefit.

D. Family and Clinical Context

Pediatric care commonly involves caregivers and multiple professionals. Intelligent systems must therefore support a network of human decision-makers rather than assuming that the robot or AI system is the sole source of rehabilitation intelligence.

E. Safety and Human Oversight

The physical interaction between robots and children increases the importance of safety assurance. Therapists need to have the ability to interpret, adjust, or override recommendations from the system. AI-generated recommendations should be within clinically defined limits.

XII. Critical Synthesis and Research Gaps

The literature shows a clear technological trajectory, but the maturity of each stage is uneven. Robotic assistance, repetitive task-oriented training, and quantitative assessment have established themselves comparatively. AI-based assessment and adaptive control are increasingly mature, but external validation and clinical generalization remain limitations. Multi-robot learning, edge intelligence, federated learning, and explainable distributed intelligence are emerging rather than being clinically established.

The most important observation is therefore that the field has algorithms. It is that the field lacks sufficiently validated integration. Individual research streams have advanced separately, while the evidence for connecting them into a longitudinal, pediatric, clinically supervised ecosystem remains limited.

Recent pediatric reviews reinforce this conclusion. AI has been applied to assessment, prediction, robotics, and assistive technologies, but the literature continues to show small samples, heterogeneous methods, limited individualized goal-setting, and limited external validation [5], [6].

Table II. Technology–evidence–research Gap Synthesis

Technology stream	Current contribution	Principal unresolved issue
Robotic assistance	Repetition, assistance, measurable training	Long-term functional benefit and generalization
AI assessment	Classification, prediction, automated quantification	External validation and clinical integration
Adaptive control	Performance-responsive assistance	Safe and clinically meaningful personalization
Multimodal systems	Complementary observations	Interoperability and shared patient representation
Multi-robot intelligence	Distributed sensing and action	Coordination and cross-device validation
Edge AI	Low-latency local inference	Resource constraints and safety assurance
Federated learning	Collaborative learning without raw-data pooling	Pediatric heterogeneity and security
Explainable AI	Interpretability and decision support	Clinician usefulness and workflow impact

A. Gap 1 — Limited Pediatric-Specific Distributed Intelligence

Distributed learning has been demonstrated in healthcare and rehabilitation robot networks, but pediatric-specific evidence remains limited. Developmental stage, diagnosis, severity, and changing therapeutic goals must be represented rather than treated as noise.

B. Gap 2 — Fragmented Multi-Robot Intelligence

Gait robots, upper-limb robots, wearable sensors, cameras, and socially assistive systems observe complementary aspects of rehabilitation but rarely contribute to a common patient-state representation.

C. Gap 3 — Incomplete Integration of Privacy and Explainability

Federated learning and explainable AI address different dimensions of trustworthy intelligence. Their combined, rehabilitation-specific evaluation remains insufficiently established.

D. Gap 4 — Limited Edge Intelligence

With Clinical Oversight Edge processing can enable fast local decisions but more solid evidence is needed to establish its integration with distributed learning, explainability and therapist supervision.

E. Gap 5 — Technical Performance Does Not Guarantee Clinical Benefit

Accuracy, latency, and kinematic improvements are important engineering measures, but they do not automatically establish better participation, adherence, independence, quality of life, or caregiver experience.

F. Gap 6 — Limited Translational Continuity

Many studies remain small, single-site, prototype-oriented, or short-term. A stronger evidence pathway is needed from laboratory validation to clinical feasibility, multicenter external validation, prospective effectiveness, and longitudinal implementation.

XIII. Review-Derived Research Direction and Future Roadmap

The gaps above suggest a research direction without implying that a complete solution has already been established. Rehabilitation intelligence may need to move from isolated devices toward distributed environments in which multiple robotic and sensing systems contribute complementary information while learning collaboratively under privacy and clinical constraints.

For pediatric rehabilitation, such a direction would require a child-centered representation, interoperable multimodal data, edge processing for time-sensitive tasks, privacy-preserving collaborative learning, clinically useful explanations, and explicit human oversight. These are research requirements to be evaluated, not assumptions that technology integration automatically improves care.

A. Near-Term Priorities

The immediate priority should be stronger evidence rather than greater algorithmic complexity. Multicenter pediatric datasets, standardized functional outcomes, external validation, transparent reporting, and pediatric-specific benchmarks are essential.

B. Medium-Term Priorities

Research can then evaluate multimodal temporal models, safe adaptive control, personalized or clustered federated learning, edge inference, and clinician-centered XAI under realistic data heterogeneity and safety constraints.

C. Long-Term Priorities

Longer-term work may investigate coordinated rehabilitation environments in which robots, wearables, clinical systems, and human professionals share meaningful information. Foundation-model-based robot learning and digital twins may contribute to this direction, but both remain emerging technologies requiring domain-specific validation.

D. Human-Centered Validation

Future trials should evaluate clinician comprehension of AI outputs, override behavior, child engagement, caregiver burden, adherence, participation, and quality of life alongside technical performance.

E. Translational Checkpoints

A responsible pathway for AI development includes laboratory validation, human–robot feasibility studies, small clinical feasibility trials, multicenter external validation, prospective comparative effectiveness research, and longitudinal implementation. The evidence burden should increase as AI moves closer to physical robotic action.

XIV. CONCLUSION

Rehabilitation technology has progressed through a sequence of responses to the limitations of preceding systems. Mechanical

assistance improved repetition and consistency; robotics combined assistance with control; sensors made rehabilitation measurable; AI made complex measurements interpretable; adaptive systems made treatment responsive; multimodal and multi-robot systems expanded the available information; edge intelligence addressed latency; and distributed learning introduced collaborative intelligence without requiring all raw data to be centralized.

The literature does not yet indicate that these technologies have converged into a mature clinical ecosystem. Instead, it reveals a transition point in which individual components are increasingly capable but their integration remains fragmented. The most important unresolved issues concern pediatric-specific generalization, genuine personalization, interoperability, privacy, explainability, safety, and longitudinal clinical effectiveness.

The pediatric context makes these issues more consequential because children grow, develop, engage differently with therapy, and receive care through networks of clinicians and caregivers. Consequently, we should not primarily judge the next stage of rehabilitation intelligence by the sophistication of its algorithms or robots.

The stronger benchmark is whether the resulting system helps clinicians understand change, helps children engage with meaningful therapy, protects sensitive information, adapts appropriately to individual needs, and produces sustained functional benefit. The literature therefore supports cautious progression toward distributed, privacy-conscious, explainable, edge-enabled, and clinically supervised rehabilitation intelligence while emphasizing that its clinical value must be demonstrated through rigorous prospective evidence.

REFERENCES

- [1] G. Nicora et al., "Systematic review of AI/ML applications in multi-domain robotic rehabilitation: Trends, gaps, and future directions," *J. Neuroeng. Rehabil.*, vol. 22, Art. no. 79, 2025, doi: 10.1186/s12984-025-01605-z.
- [2] A. Gonzalez, L. Garcia, J. Kilby, and P. McNair, "Robotic devices for paediatric rehabilitation: A review of design features," *BioMed. Eng. OnLine*, vol. 20, Art. no. 89, 2021, doi: 10.1186/s12938-021-00920-5.
- [3] R. Retnaningsih et al., "Robotics applications in neurology: A review of recent advancements and future directions," 2023.
- [4] N. Al-Rahmani et al., "The past, present and future of control architectures in lower-limb cable-driven robots for gait rehabilitation," *J. Neuroeng. Rehabil.*, vol. 23, Art. no. 188, 2026, doi: 10.1186/s12984-026-01982-z.
- [5] V. C. Kaelin et al., "Artificial intelligence in rehabilitation targeting the participation of children and youth with disabilities: Scoping review," *J. Med. Internet Res.*, vol. 23, no. 11, Art. no. e25745, 2021, doi: 10.2196/25745.
- [6] N. Naghdi et al., "Use of artificial intelligence for outcome assessment in pediatric rehabilitation: A scoping review," *J. Neuroeng. Rehabil.*, vol. 23, Art. no. 123, 2026, doi: 10.1186/s12984-026-01927-6.
- [7] D. F. Brown and S. Q. Xie, "Effectiveness of intelligent control strategies in robot-assisted rehabilitation—A systematic review," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 32, pp. 1828–1840, 2024, doi: 10.1109/TNSRE.2024.3396065.
- [8] A. Zhaisanbek et al., "Social robots in cognitive and speech rehabilitation for children with cerebral palsy: A scoping review," *J. Neuroeng. Rehabil.*, vol. 23, Art. no. 34, 2026, doi: 10.1186/s12984-025-01852-0.
- [9] J. Li, Q. Zhu, M. N. Mohd Noor, and C. C. Foong, "Effectiveness and usability of artificial intelligence-powered assistive technologies in supporting daily activities of children with cerebral palsy: A systematic review," *Ann. Med.*, 2026, doi: 10.1080/07853890.2026.2649427.
- [10] X. Jiang et al., "Federated joint learning of robot networks in stroke rehabilitation," in *Proc. 33rd IEEE Int. Conf. Robot Human Interactive Commun. (RO-MAN)*, 2024, pp. 1294–1300, doi: 10.1109/RO-MAN60168.2024.10731407.
- [11] J. Xu et al., "Federated learning for healthcare informatics," *J. Healthc. Inform. Res.*, 2021, doi: 10.1007/s41666-020-00082-4.
- [12] T. Tunduny and B. Shibwabo, "Explainable AI approaches in federated learning: Systematic review," *JMIR AI*, vol. 5, e69985, 2026, doi: 10.2196/69985.
- [13] A. Chaddad et al., "Explainable, domain-adaptive, and federated artificial intelligence in medicine," *IEEE J. Biomed. Health Inform.*, 2022.
- [14] S. Liu et al., "Leveraging explainable artificial intelligence to optimize clinical decision support," *J. Amer. Med. Inform. Assoc.*, vol. 31, no. 4, pp. 968–974, 2024, doi: 10.1093/jamia/ocae019.
- [15] S. Ghanvatkar and V. Rajan, "Evaluating explanations from AI algorithms for clinical decision-making: A social science-based approach," *IEEE J. Biomed. Health Inform.*, vol. 28, no. 7, pp. 4269–4280, 2024, doi: 10.1109/JBHI.2024.3393719.
- [16] K. Dasaradharami Reddy and T. R. Gadekallu, "A comprehensive survey on federated learning techniques for healthcare informatics," *Comput. Intell. Neurosci.*, 2023, Art. no. 8393990, doi: 10.1155/2023/8393990.
- [17] X. Xiao et al., "Robot learning in the era of foundation models: A survey," *Neurocomputing*, vol. 638, Art. no. 129963, 2025, doi: 10.1016/j.neucom.2025.129963.
- [18] I. D. Mienye, E. Ezenogho, and C. Modisane, "Deep reinforcement learning in the era of foundation models: A survey," *Computers*, vol. 15, no. 1, Art. no. 40, 2026.
- [19] Q. Sun, A. Akman, and B. W. Schuller, "Explainable artificial intelligence for medical applications: A review," 2024.
- [20] "Digital twin-driven intelligent healthcare systems," review, 2024.