

Human-AI Cooperation in Healthcare and Rehabilitation

Austin Brockmeier, PhD;¹ Panagiotis Artemiadis, PhD;¹ Hacene Boukari, PhD;² Chris Callison-Burch, PhD;³ Eric Eaton, PhD;³ Joel Harley, PhD;⁴ Thomas Powers, PhD;¹ Jose C. Principe, PhD;⁴ Darcy Reisman, PhD;¹ Cathy Wu, PhD¹

1. University of Delaware
2. Delaware State University
3. University of Pennsylvania
4. University of Florida

Abstract

Rehabilitation after injury or to manage chronic health conditions requires continuous reassessment and intervention across time scales ranging from seconds to months. Advances in sensors and data collection, coupled with new technology to administer interventions, create numerous possibilities—including at-home care. The increased capabilities enable automated analysis and control using artificial intelligence (AI). In this essay, we analyze the need, the potential and the requirements for an intense and enduring physical human-AI cooperation framework, i.e., a symbiosis, where both AI and humans contribute to realize improved solutions. The focus is the development of knowledge and expertise to realize a new generation of AI-enabled therapy for the next decades. With an aging population, prevalence of stroke and chronic diseases, there is a demand for more efficient and effective rehabilitation powered by human-AI cooperation, especially in cases that enable remote participation in areas with limited access. This essay analyzes how the potential for advances in human-AI cooperation can impact rehabilitation in Delaware.

Introduction

This paper discusses the state-of-the-art and the potential, both near-term and long-term, to improve therapy and rehabilitation through human-AI cooperation, along with initial plans as developed by an interdisciplinary team of researchers. The use case of personalized rehabilitation and healthcare is intrinsically challenging, as it requires decision making with incomplete, multimodal data at several time scales about a patient's health informed by experience and knowledge of current evidence. In particular, stroke rehabilitation involves efficiently gathering information through diverse methods¹ on motor-sensory deficits, given the patient's previous abilities, to outline exercises and assistive robotics, which will most effectively re-establish lost abilities—in an environment limited by possible impairments in speech and cognition.² Both the patient's health and the clinician's understanding of the patient's current health and goal state change through multiple sessions due to the treatment and unobservable factors, with the vast majority of time being outside of clinical observation. In this context, advances in artificial intelligence technologies can enable the seamless analysis and integration of static and dynamic patient data collected by multimodal sensors within and outside of the clinic.^{3–5} This includes real-time closed-loop sensing with robotic,⁶ computer, and machine interfaces for patient sensory stimulation and monitoring, expert-guided (human-in-the-loop) contextualization of patients based on relevant information into strata for optimized

treatment, and personalized AI-enabled communication interfaces for expanding patient-clinician dialogue for patient support and engaging in out-of-clinic exercises and assessment.

Sensorimotor disabilities resulting from neurological disorders or injuries are a pressing challenge⁷ as rehabilitation is both financially burdensome and labor-intensive. Prevailing rehabilitation practices lack customization. The goal is to harness technological advancements in robotics and sensing to tailor treatment to each patient's characteristics by developing AI models proficient in leveraging limited patient-specific data and existing knowledge to project potential functional outcomes of interventions into the future. The data sources encompass patient interactions with clinicians, robotic mechanisms in clinical environments, and wearable sensors capturing patient activity in daily life settings.

Challenges and Potential

Although progress has been made based on emerging AI techniques, current AI cannot yet truly act as a partner in decision making due to limitations in its foundations and the realities of the use case. For example, machine learning approaches for data processing do not provide conceptual grounding to reason about the information content of heterogeneous data, which is necessary to communicate uncertainty and optimize the gathering of additional data. AI must consider a patient's health as a partially observable state and clinicians need to be able to specify goals and constraints directly. The validity and relevance of data and knowledge are continually evolving at multiple scales beyond the dynamics of individual patients; new sensors or treatments will become available while others become obsolete; new evidence will lead to changes in standards of care; and the population of patients will change. Modeling and/or optimizing treatment in such nonlinear, time-evolving systems is a challenge. Even with improved sensing, it is not trivial to extract meaningful information from data to identify the causal relationships of variables of interest in a dynamic environment. Additionally, there is no blueprint on how to divide the roles between human experts and AI to maximize human-AI performance for clinical decision making.

AI for rehabilitation has unique specialized algorithms and ethical implications. It requires creating a research nexus between AI researchers, physical therapy researchers, community stakeholders, and industry partners ranging from computing hardware/software and medical equipment manufacturers to healthcare providers. Initially, pilot projects operating at different scales are needed to validate developments in real-world settings and identify new challenges and solutions as foundational contributions are realized. Patient and clinician/physical therapist feedback should be sought during *co-design* to assess how user-friendly and efficient these AI-assisted systems are in a real-world therapeutic context. There is also a need to connect and educate a spectrum of people from AI researchers to health-care practitioners.


In current rehabilitation research, models of human physiology and psychology blend mechanistic and empirical models at different scales of accuracy. Yet, even with this imperfect understanding using limited data, clinicians characterize subjects and personalize treatment, so their role in guiding the final architecture choices is fundamental.⁸ At the same time, clinicians stand to benefit from AI's computational power to integrate massive and diverse information in data repositories, from natural language to information-rich signals from video and sensors and to identify subtle patterns in massive quantitative data. AI is also advantageous because its continuous operation does not lead to degradation of performance due to the stress or

exhaustion as happens with humans. However, novel human-AI systems capable of leveraging the intelligence of both the human and AI in cooperative feedback through *grounded concepts* is needed.

In rehabilitation and healthcare, AI operation must be specific to the characteristics of an individual, which evolve at a variety of timescales. The AI systems must be *trustworthy*, *instructable* and *aligned* with the patient's goals as supervised by clinicians. The rehabilitation of sensorimotor disabilities through physical therapy involves human motor learning, which relies on the nervous system's ability to integrate sensory information, control movement, and plasticity, governed by an individual's intrinsic reward. These characteristics are desired in AI systems. Thus, an improved understanding of natural learning is synergistic with advances in human-AI cooperation.

Vision for AI-enabled Human-AI Cooperation. We imagine an AI-enabled system for human-AI cooperation (Figure 1) that can dialogue with the human expert to inform if a particular intervention could maximize the rehabilitation goals for a given patient. This will improve the outcomes through the clinician's (therapist's) enhanced abilities to make sense of a variety of sensors and data and the patients' ability to interact with AI-enabled devices to support their participation in therapy. In addition to sensors and data, the AI system can directly use instructions and feedback from users (both therapists and patients) and monitor (with informed consent) the patient and patient-therapist interaction via multiple modalities. Even with these data sources, the AI's information may be limited compared to the users. To ensure that the therapist's mental model of the patient aligns with the AI's characterization, we propose to use a digital twin and simulation environment to provide an avatar in a virtual embodiment for exploring and displaying structural and functional aspects of the AI's model of a patient.⁹ This can serve multiple purposes including providing a patient with a guide for exercises and visualization of future health states.

Figure 1. Illustration of the Human-AI Cooperation in the Context of Rehabilitation.




This requires modeling complex, stochastic, and dynamic processes from heterogeneous data at multiple scales and for estimating and adapting optimal policies for specific individuals, contexts, or strata of human populations in the presence of new information and direct human instruction. There is a need for cross-cutting approaches for extracting relevant information from data, fusing real and simulated data—while robustly handling distribution shifts, and

creating accessible and multimodal human-AI interfaces that ensure alignment of AI operation with user intent through direct instruction and goal specification grounded in human language (verbal and non-verbal forms, data exemplars, and other knowledge representations). The grounded interfaces will also allow AI explanations of model predictions and decision-making, ensuring alignment for the human-AI cooperation during both model adaptation and real-time operation. AI-enabled robotic agents can be used for rehabilitation and assistance, including applications in rehabilitation exercise robots, prosthetics, and orthotics. More comprehensively, human-AI cooperation requires ethical framework and neuroscientific perspectives of, including studies of the neural correlates of trust, error, and reward, using brain-computer interfaces.

We envision AI as a tool to make sense of the revolutionary technical capabilities that every day extend the quantitative assessment of physical health both inside and outside of the clinic to provide real-time interaction for analysis. Using the increased sensing, we envision AI empowering adaptive mechanisms in rehabilitation (including robotic assistance and functional stimulation) to provide personalized and effective interventions for patients. In particular, robot-assisted interventions could be customized based on the predictive abilities of the AI models. This approach allows the system to determine the most suitable intervention for each patient based on their unique functional needs. Similarly, using advances in natural language interfaces and augmented/virtual reality, AI-based virtual interactions can provide virtual coaching and feedback that is critical for optimal rehabilitation outcomes which require specific activity and guided, individualized exercises outside of the clinical setting. The physical therapist or clinician and the patient would use the *digital twin* to review the model's proposed approach, refine it by instruction, and the AI would personalize it to the patient and clinician through *preference elicitation*. At home, the digital twin, wearable sensors, and natural language interfaces would enable an interactive and adaptable session. This type of personalized care can enhance compliance with exercise instruction and improve outcomes.¹⁰ The physical therapist can interact with the AI to modify the exercises, prompts, and instructions at the next clinic visit. Such human-in-the-loop optimization can result in more responsive and efficient care that represents a critical next step in rehabilitation and healthcare that has been sorely lacking.

Our description of a system for human-AI cooperation (Figure 2) is meant to convey the overall information flow and processing and interfaces involved. Loops will allow human-AI co-exploration, policy improvement, user preference elicitation, and AI-guided data exploration.

Figure 2. Information Flow in an AI System for Human-AI Cooperation



Potential Milestones for Human-AI Cooperation in Rehabilitation

Goal 1. Provide rapid analysis of information-rich synchronized multimodal measurements (video, motion and force capture, brain activity) of cognitive state, muscle activity, and sensory-motor feedback to help characterize a patient's state and progress in recovery towards goals. This requires the following:

- Integration of measurements of varying fidelity outside of the clinic with those inside the clinic.
- Identification of relevant patterns in multimodal data—both static patterns and longitudinal trends.
- Instantiation of digital twins for interactive exploration and simulation of patient activity.
- Characterization of a patient as being of a particular group or strata relevant for treatment.
- Retrieval of relevant cross-sectional data and knowledge to support statistical reasoning.
- Uncertainty quantification of predictions to inform subsequent measurements.

Goal 2. Model the longitudinal dynamics of patient health trajectories to aid in diagnosis and treatment. This requires the following:

- Simulation of patient rehabilitation with a digital twin to explore how function and activities are predicted to be impacted through different treatment to aid decision making.
- Progression tracking with automatically updating data products and early detection of changes.
- Expert-in-the-loop algorithms for optimizing patient-specific policies reactive to real-time data.
- Calibration of closed-loop systems (robotics and medical devices) for AI-enabled therapy.

These milestones are ambitious and go well beyond the current state of practice. Sensors outside of the clinic will support wider access via telehealth and interactive and aligned AI-enabled systems will support remote therapy outside of visits. Long term, a universal model for patient health/rehabilitation trajectories at scale grounded in knowledge-based concepts will enable

human-AI cooperation for generating hypotheses and insights from studies that currently are not comparable due to differences in interventions and data collection methods. The use of AI will provide a major advance in our ability to interpret data efficiently and precisely, to gain insight on the underlying deficits and predict progress through rehabilitation, and to achieve AI-aided clinical decision-making.

Key Components

Integration

Grounded AI requires an ability to make sense of varied data. Most clinical contexts, and indeed human decision making, naturally involve multimodal data. Consequently, the integration of diverse multimodal data streams is critical for advancing our understanding and computational abilities, particularly in the context of enhancing human-AI collaboration. Multimodal data, which may include video, motion capture, audio, text, streaming sensor outputs, and human-provided instructions, presents unique challenges due to its heterogeneous nature. These data types exist across various spatiotemporal scales. Frequently the sensory data is incomplete or imprecise, requiring sophisticated frameworks that can not only handle this diversity but also leverage it to empower human decision-making alongside AI systems.

There is a need for a framework for the effective integration of such multimodal information, facilitating abstracted learning and reasoning that operate across multiple scales. This is essential for grounding complex data with inherent “missingness”, extracting meaningful information from them in a manner that supports robust learning and inference by both humans and AI systems. Critical capabilities for reliable human-AI decision-making in unpredictable environments include uncertainty quantification and bounds on performance in the presence of changing data distributions.

Data Sources

The top-ranked University of Delaware (UD) Physical Therapy program can provide data from multi-site randomized clinical trials in rehabilitation. The program has generated large and rich research data, including human motion capture, force and muscle activity, and electronic health record (EHR), to address complex problems of activity and participation.^{1,11–15} Another data source is deidentified multimodal data of primary care clinical encounters collected by Penn Engineering and Medicine, including video, audio, transcripts, EHR, and audit logs.¹⁶ The data provides an unprecedented view in an ML-ready format while preserving patient privacy. The long-term goal is to enable in-clinic AI analysis of a patient's data to provide context of a patient's physical, cognitive, and emotional health state, support real-time clinical decision-making, alert clinicians to relevant patterns and changes, and improve the clinical processes to achieve a higher standard of care.

Grounded AI for human-AI dialogue and explanations. Next-generation AI must be able to receive direct instructions from users in various modalities, and it must be able to provide explanations that are both understandable to users and faithful to the AI reasoning processes.^{17,18} AI-human interfaces should be bi-directional: in the human-to-AI direction, instructions and other interactions will form the context of the state representation that is used in the AI's subsequent predictions; and in the AI-to-human direction, the model needs to be able to provide

explanations of its predictions or outputs.¹⁹ The first direction requires understanding the best way to represent these instructions as part of the state, objective, or constraints on the policies or actions. The second direction builds on our work on explainable AI.¹⁹ Instructable multimodal language models can provide explanations of their predictions, essential in healthcare and other safety critical applications. In particular, large language models (LLMs) trained on extremely large bodies of text and dialogues have made interactions between machines and humans seem deceptively solved. LLMs fine-tuned on instructions have demonstrated the ability to follow instructions and engage in open-ended dialogue. Building on this, ChatGPT and other dialogue systems²⁰ leverage instruction fine tuning on a large corpus of conversational data to enable engaging in natural back-and-forth conversations while following instructions and context. But human interaction is not just limited to text; it is holistic and multimodal and may need to be personalized. Often, humans interact via gestures and body language as well as by conversational common ground that includes shared visual referents. Large vision-language models (LVLMs) that integrate visual inputs along with text provide an opportunity to study communication at the intersection of vision and text. Can these techniques for creating a shared representation between language and vision be extended to other modalities like human poses²¹ gathered from sensor data? For instance, could an AI agent helping a patient with their physical therapy exercises be able to view the patient's motions and give instructions on what other motions to make using understandable language?

To deliver multimodal instructability and faithful explanations that mitigate hallucinations, we envision novel approaches for verbally describing relationships and patterns in sensor data, visualizations, knowledge graphs and other media, and then training networks to embed these in shared embedding spaces. One potential approach is to augment existing non-language data with language descriptions using LLMs and LVLMs. We propose to extend our framework for synthetic data generation and multimodal model training via complex prompting workflows.²² We envision this synthetic data may take the form of language that describes data signals, including producing chain-of-thought-style reasoning²³ and instruction-style data,^{24,25} and then training multimodal generative models on this data. By building frameworks for synthetic data generation and multimodal model training, we will be able to augment data collected in other portions of our AI2HAI program to create specialized AI systems built for our use cases including physical therapy, and that will provide adaptive high quality grounded explanations to improve AI-patient communication.

Personalization

Human-centered decision-making requires personalized, resource-efficient, and equitable solutions. We consider a framework with a single statistical model for a given task where the human-centric data is incorporated into the state or influences the objective, actions, or constraints. Framed in the formalisms of sequential decision making, a belief state distills relevant information from a history of observations, to choose actions based on a policy that is optimal in maximizing reward for each specific subject. Such a policy is personalized by the context provided by information about the individual gathered from data integration or from human instruction. Instructability dictates that as additional human input is given, any policy is flexible such that a modified state based on the observed instructions ensures actions that satisfy these instructions. Once a state representation is informed with human feedback and personalized data, it provides the basis for optimizing policies. Human interpretation of

uncertainty quantification in the state is essential. Users would benefit from knowing how trustworthy the machine prediction is, especially in health care.


The importance of digital twins in personalization. Digital twins couple computational models with a physical counterpart^{26–28} that can be dynamically updated through bidirectional data flow as conditions change.²⁹ Biophysically based models have the potential to improve biomedical decision-making at the individual level, but generating models capable of directly informing patient treatments remains a significant challenge,³⁰ as generic models are insufficient for characterizing population diversity,^{31,32} necessitating personalization. Personalization requires measurements of the parameters underlying a model through batteries of sensor-based data collections, which can be exceedingly difficult and costly.³³ This motivates the improvement of the efficiency and precision of parameter determination to ensure alignment at the level of an individual. Ideally, one could decode (solve the inverse problem) of identifying the digital twin's parameters from an individual's state representation, itself computed from integrative multimodal processing including self-supervised learning, which will greatly improve the quality of personalization by capturing the relevant information about a given individual. Then, the digital twin can enable subsequent simulation from the individual's state. The digital twin-based simulations themselves will create virtual data that can be fed back to the integrative multimodal data processing and used to compare to the original state presentation and decoded parameters, with cyclic loss functions that penalize the differences used to update decoder and in updates of the integrative multimodal processing. Finally, models capable of capturing the long-term dynamics of the state can be used with the digital twin to simulate future states.

The current state-of-the-art for digital twins include NVIDIA's OmniVerse, which bridges collected data and AI systems with the 3D world building. This can link video/motion capture and sensor measurements with the digital twin. OnmiVerse (and by extension, IsaacSim and IsaacGym) provides a common platform that has the capabilities to intake data and models and seamlessly and efficiently *integrate reinforcement learning and other AI algorithms with physics engines*. IsaacSim provides a robust physics engine for model patients in rehabilitation and IsaacGym provides tools for seamless integration with reinforcement learning, enabling optimization of control algorithms for various robotics platforms, as in our previous work.³⁴ To be realistic it will require mechanistic models or data from these mechanistic models (e.g., OpenSIM,³⁵ MuJoCo³⁶) that are regularly used to model biomechanics and rehabilitation processes. These widely tested and validated models can be used to create the initial digital twins.

Human-AI Symbiosis

Alignment of AI decision making can use hybrid RL that merges offline data (collected from human behavior) with online simulations involving digital twins to provide a means to produce optimal policies that are supported by human behavior. When users exploit AI-aided decision making, their new behavior is a step toward optimization, and their new behavior can be used in subsequent AI policy optimization. Through iterations, this will enable human-in-the-loop co-evolution of policies that can safely explore increasingly complicated decision spaces—achieving unprecedented optimality as shown in Figure 3. We envision human-AI co-evolution to explore the policy space efficiently and safely, while maximizing measurable and functional outcomes.

Figure 3. Evolution of AI in Rehabilitation



Establishing a trustworthy human-AI team is crucial for human-centric decision-making and actions, especially in rehabilitation with clinical intervention and robotic assistance. There is a need for cross-cutting approaches to measure human trust at behavioral and neural levels, where trust in human-autonomy interaction is defined in situations characterized by uncertainty,³⁷ based on extracting levels of human trust in AI-driven agents through real-time monitoring of neural correlates.³⁸ In particular, EEG may be used to quantify in real-time the effectiveness of human-AI teaming.³⁸ This real-time estimation of trust is a pivotal factor in guiding the AI system to align with the human teammate, revolutionizing the effectiveness of human-AI cooperation. This can be extended to cases where trust has to be repaired.³⁹ When trust is established, we expect Human-AI symbiosis to emerge as the AI becomes better adapted to meet the needs and as communication modes of human collaborators and the human collaborators better understand how to communicate and benefit from AI capabilities.

While improved sensing can provide functional and structural information, and patient questionnaires can provide understanding of activities and participation, the ability to understand how a patient feels or discomfort is poorly quantified. We envision using natural language-based interactive patient dialogues with patients that build trust and a better understanding of patient perspectives. Long term, we envision deriving neural correlates of reward to directly guide the AI adaptation, which was proposed by our team for brain-machine interfaces⁴⁰ and has been improved steadily in animal models.⁴¹ Brain machine interfaces provide a motivating use case of human-AI symbiosis in rehabilitation, where the machine directly senses and decodes the brain's activity regarding musculoskeletal control and reward to control external robotics (exoskeletons).⁴² A key lesson from our research⁴⁰ is that the brain of each subject is not only unique but changes day to day,⁴³ one of the challenges and potentials of aiding in neuromuscular rehabilitation. Building on our work in stroke rehabilitation,⁴⁴ there is the potential to combine EEG-based correlates of reward with muscle activity in closed-loop robotics and a virtual avatar to induce neuromuscular synchrony to speed recovery.

Predictive modeling for AI-powered robot-assisted personalized interventions. Our prior research demonstrated that complex neuromusculoskeletal models could simulate robot-assisted therapies for both healthy individuals^{45,46} and those with mobility impairments.⁴⁷ These models effectively replicate healthy and impaired walking with high accuracy, yet they fall short in capturing motor learning and adaptation—key elements in the rehabilitation process. AI may be

able to predict how patients adapt to specific robot-assisted interventions, by modeling the process of motor learning at the brain level. This breakthrough has the potential to revolutionize sensorimotor rehabilitation by introducing a new generation of modeling frameworks. This approach allows therapy interventions to be simulated, optimized, and personalized before being applied during treatment, allowing for a more targeted and effective therapy experience. The ability to model and thus simulate rehabilitation will be essential for further enabling personalized decision making and digital twinning. A range of wearable technology and robot-assisted rehabilitation devices can be integrated into the simulation environment. Doing so allows us to fine-tune various parameters—such as intervention type, frequency, and therapy duration—to achieve better functional outcomes. The ultimate goal is to create interventions that effectively promote motor adaptation, improving each individual's rehabilitation.

Accessible Natural Language (NL) Interfaces. Human-AI partnering means that there must be interaction between the machine and the human that is understandable and grounded for both the AI and the human. The AI interfaces must themselves be adaptable and tailored to interact with clinicians (for AI partnering) and end-users (patients). A participatory user-centered design practices should be adopted involve practitioners, their clients, and other stakeholders throughout the design process, starting from involving diverse stakeholders in the requirements gathering and using ethnographic observations, contextual inquiry, group and individual interviews, and having them interact with a series of iterative prototypes leading to validation in randomized control trials. Interacting with the patients requires extra care in safety.⁴⁸ In addition, patients in our studies may experience aphasia or other cognitive impairments which must be accounted for in the participatory design process—previous projects began to tackle these challenges.^{49–52} Research on health and rehabilitation has a particular onus to reach vulnerable and underserved populations due to significant health disparities and disparities in social determinants,⁵³ such as education, employment, and socio-economic status. People with disabilities are one such vulnerable group.⁵⁴ To ensure accessibility, communications generated by AI systems should be personalized through multiple communication modalities to achieve the concept of “born accessible”.⁵⁵ Accessibility comes in many forms including being able to translate between graphics and text taking into account characteristics of the reader,⁵⁶ and text simplification. It could also mean adapting suggestions to users who may need different support or may respond to different therapeutic exercises with varying degrees of enthusiasm⁵⁷ or who might engage more or less with certain types of interfaces—including the choice of AI personalities.⁵⁸

AI ethics framework. There is a threat of both under- and over-estimation of the risks of AI, so the study of AI ethics must weigh risks and prioritize mitigation strategies. An ethics framework for rehabilitation and physical therapy integrates standard-of-care guardrails for treatment, e.g., patient consent-to-treat and maintenance of patient-therapist trust. The ethics framework will need to respond to functional gains of AI by supplying programmed constraints⁵⁹ in human-AI cooperation applications, such as in AI-enabled personalized healthcare interventions.³⁹

Conclusion

With human-AI symbiosis, new possibilities for optimizing rehabilitation interventions are supported as the continual adaptation of empirical and biophysical models and expert behavior over time gather supporting evidence at the cutting edge. Together, this will support human-AI

co-evolution of new policies—accelerating the exploration of possibilities while still guided by human expertise. Creating AI systems that enable these transformative capabilities will require many components and diverse expertise as outlined in this work.

The authors may be contacted at wuc@udel.edu.

References

1. Miller, A. E., Russell, E., Reisman, D. S., Kim, H. E., & Dinh, V. (2022, June 17). A machine learning approach to identifying important features for achieving step thresholds in individuals with chronic stroke. *PLoS One*, 17(6), e0270105. <https://doi.org/10.1371/journal.pone.0270105> PubMed
2. French, M. A., Cohen, M. L., Pohlig, R. T., & Reisman, D. S. (2021, May). Fluid cognitive abilities are important for learning and retention of a new, explicitly learned walking pattern in individuals after stroke. *Neurorehabilitation and Neural Repair*, 35(5), 419–430. <https://doi.org/10.1177/15459683211001025> PubMed
3. Blanton, S., Cotsonis, G., Brennan, K., Song, R., Zajac-Cox, L., Caston, S., . . . Kesar, T. (2023, November 24). Evaluation of a carepartner-integrated telehealth gait rehabilitation program for persons with stroke: Study protocol for a feasibility study. *Pilot and Feasibility Studies*, 9(1), 192. <https://doi.org/10.1186/s40814-023-01411-1> PubMed
4. Silva-Batista, C., Wilhelm, J. L., Scanlan, K. T., Stojak, M., Carlson-Kuhta, P., Chen, S., . . . King, L. A. (2023, October 13). Balance telerehabilitation and wearable technology for people with Parkinson’s disease (TelePD trial). *BMC Neurology*, 23(1), 368. <https://doi.org/10.1186/s12883-023-03403-3> PubMed
5. Miller, A., Collier, Z., & Reisman, D. S. (2022, October 14). Beyond steps per day: Other measures of real-world walking after stroke related to cardiovascular risk. *Journal of Neuroengineering and Rehabilitation*, 19(1), 111. <https://doi.org/10.1186/s12984-022-01091-7> PubMed
6. Srivastava, S., Kao, P. C., Reisman, D. S., Scholz, J. P., Agrawal, S. K., & Higginson, J. S. (2016, October). Robotic assist-as-needed as an alternative to therapist-assisted gait rehabilitation. *International Journal of Physical Medicine & Rehabilitation*, 4(5), 370. <https://doi.org/10.4172/2329-9096.1000370> PubMed
7. World Health Organization. (n.d.). *Neurological disorders: public health challenges*. Retrieved from <https://www.who.int/publications/i/item/9789241563369>
8. Miller, A., Pohlig, R. T., Wright, T., Kim, H. E., & Reisman, D. S. (2021, October). Beyond physical capacity: Factors associated with real-world walking activity after stroke. *Archives of Physical Medicine and Rehabilitation*, 102(10), 1880–1887.e1. <https://doi.org/10.1016/j.apmr.2021.03.023> PubMed
9. Seth, A., Hicks, J. L., Uchida, T. K., Habib, A., Dembia, C. L., Dunne, J. J., . . . Delp, S. L. (2018, July 26). OpenSim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PLoS Computational Biology*, 14(7), e1006223. <https://doi.org/10.1371/journal.pcbi.1006223> PubMed

10. Davergne, T., Meidinger, P., Dechartres, A., & Gossec, L. (2023, July 13). The effectiveness of digital apps providing personalized exercise videos: Systematic review with meta-analysis. *Journal of Medical Internet Research*, 25(1), e45207. <https://doi.org/10.2196/45207> PubMed
11. French, M. A., Daley, K., Lavezza, A., Roemmich, R. T., Wegener, S. T., Raghavan, P., & Celnik, P. (2023). A learning health system infrastructure for precision rehabilitation after stroke. *American Journal of Physical Medicine & Rehabilitation*, 102(2S Suppl 1), S56–S60. <https://doi.org/https://doi.org/10.1097/PHM.0000000000002138>
12. Miller, A., Pohlig, R. T., & Reisman, D. S. (2022, August). Relationships among environmental variables, physical capacity, balance self-efficacy, and real-world walking activity post-stroke. *Neurorehabilitation and Neural Repair*, 36(8), 535–544. <https://doi.org/10.1177/15459683221115409> PubMed
13. Harbourne, R. T., Dusing, S. C., Lobo, M. A., McCoy, S. W., Koziol, N. A., Hsu, L.-Y., . . . Sheridan, S. M. (2021, February 4). START-Play physical therapy intervention impacts motor and cognitive outcomes in infants with neuromotor disorders: A multisite randomized clinical trial. *Physical Therapy*, 101(2), pzaa232. <https://doi.org/10.1093/ptj/pzaa232> PubMed
14. Su, W.-C., Cleffi, C., Srinivasan, S., & Bhat, A. (2023, November 1). Telehealth versus face-to-face fine motor and social communication interventions for children with autism spectrum disorder: Efficacy, fidelity, acceptability, and feasibility. *The American Journal of Occupational Therapy*, 77(6), 7706205130. <https://doi.org/10.5014/ajot.2023.050282> PubMed
15. Master, H., Bley, J. A., Coronado, R. A., Robinette, P. E., White, D. K., Pennings, J. S., & Archer, K. R. (2022, February 15). Effects of physical activity interventions using wearables to improve objectively-measured and patient-reported outcomes in adults following orthopaedic surgical procedures: A systematic review. *PLoS One*, 17(2), e0263562. <https://doi.org/10.1371/journal.pone.0263562> PubMed
16. The Observer Project. (n.d.). *Welcome to the Observer Project*. Perelman School of Medicine at the University of Pennsylvania. Retrieved from <https://www.med.upenn.edu/observer/>
17. Lyu, Q., Apidianaki, M., & Callison-Burch, C. (2024). *Towards faithful model explanation in NLP: A survey* (arXiv:2209.11326). arXiv. <https://doi.org/10.48550/arXiv.2209.11326>
18. Lyu, Q., Havaldar, S., Stein, A., Zhang, L., Rao, D., Wong, E., . . . Callison-Burch, C. (2023). *Faithful chain-of-thought reasoning* (arXiv:2301.13379). arXiv. <https://doi.org/https://doi.org/10.18653/v1/2023.ijcnlp-main.20>/arXiv.2301.13379
19. Yang, Y., Panagopoulou, A., Zhou, S., Jin, D., Callison-Burch, C., & Yatskar, M. (2023). *Language in a bottle: language model guided concept bottlenecks for interpretable image classification* (arXiv:2211.11158). arXiv. <https://doi.org/https://doi.org/10.1109/CVPR52729.2023.01839>/arXiv.2211.11158

20. Thoppilan, R., Freitas, D. D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.-T., . . . Le, Q. (2022). *LaMDA: language models for dialog applications* (arXiv:2201.08239). arXiv. <https://doi.org/10.48550/arXiv.2201.08239>
21. Huang, Y., Wan, W., Yang, Y., Callison-Burch, C., Yatskar, M., & Liu, L. (2024). *CoMo: controllable motion generation through language guided pose code editing* (arXiv:2403.13900). arXiv. <https://doi.org/10.48550/arXiv.2403.13900>
22. Patel, A., Raffel, C., & Callison-Burch, C. (2024). *DataDreamer: A tool for synthetic data generation and reproducible LLM workflows* (arXiv:2402.10379). arXiv. <https://doi.org/https://doi.org/10.18653/v1/2024.acl-long.208/arXiv.2402.10379>
23. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., . . . Zhou, D. (2023). *Chain-of-thought prompting elicits reasoning in large language models* (arXiv:2201.11903). arXiv. <https://doi.org/10.48550/arXiv.2201.11903>
24. Mishra, S., Khashabi, D., Baral, C., & Hajishirzi, H. (2022). *Cross-task generalization via natural language crowdsourcing instructions* (arXiv:2104.08773). arXiv. <https://doi.org/https://doi.org/10.18653/v1/2022.acl-long.244/arXiv.2104.08773>
25. Wang, Y., Mishra, S., Alipoormolabashi, P., Kordi, Y., Mirzaei, A., Arunkumar, A., . . . Khashabi, D. (2022). *Super-natural instructions: generalization via declarative instructions on 1600+ NLP tasks* (arXiv:2204.07705). arXiv. <https://doi.org/10.48550/arXiv.2204.07705>
26. Lindbeck, E. M., Diaz, M. T., Nichols, J. A., & Harley, J. B. (2023, December). Predictions of thumb, hand, and arm muscle parameters derived using force measurements of varying complexity and neural networks. *Journal of Biomechanics*, 161, 111834. <https://doi.org/10.1016/j.jbiomech.2023.111834> PubMed
27. Diaz, M. T., Harley, J. B., & Nichols, J. A. (2024, February 1). Sensitivity analysis of upper limb musculoskeletal models during isometric and isokinetic tasks. *Journal of Biomechanical Engineering*, 146(2), 021005. <https://doi.org/10.1115/1.4064056> PubMed
28. Tappan, I., Lindbeck, E. M., Nichols, J. A., & Harley, J. B. (2024, March). Explainable AI elucidates musculoskeletal biomechanics: A case study using wrist surgeries. *Annals of Biomedical Engineering*, 52(3), 498–509. <https://doi.org/10.1007/s10439-023-03394-9> PubMed
29. National Academies. (n.d.). *Foundational research gaps and future directions for digital twins 2024*. Retrieved from <https://www.nationalacademies.org/publications/26894>
30. Katsoulakis, E., Wang, Q., Wu, H., Shahriyari, L., Fletcher, R., Liu, J., . . . Deng, J. (2024, March 22). Digital twins for health: A scoping review. *NPJ Digital Medicine*, 7(1), 77. <https://doi.org/10.1038/s41746-024-01073-0> PubMed
31. Castro, M. N., Rasmussen, J., Bai, S., & Andersen, M. S. (2019, June 11). Validation of subject-specific musculoskeletal models using the anatomical reachable 3-D workspace. *Journal of Biomechanics*, 90, 92–102. <https://doi.org/10.1016/j.jbiomech.2019.04.037> PubMed
32. Goislard De Monsabert, B., Edwards, D., Shah, D., & Kedgley, A. (2018, January). Importance of consistent datasets in musculoskeletal modelling: A study of the hand and

wrist. *Annals of Biomedical Engineering*, 46(1), 71–85. <https://doi.org/10.1007/s10439-017-1936-z> PubMed

33. Kerkhof, F. D., van Leeuwen, T., & Vereecke, E. E. (2018, November). The digital human forearm and hand. *Journal of Anatomy*, 233(5), 557–566. <https://doi.org/10.1111/joa.12877> PubMed
34. Scully, C. (2024). *Reinforcement learning-based controller for quadruped locomotion over compliant terrain* [University of Delaware]. <https://udspace.udel.edu/handle/19716/35100>
35. Delp, S. L., Anderson, F. C., Arnold, A. S., Loan, P., Habib, A., John, C. T., . . . Thelen, D. G. (2007, November). OpenSim: Open-source software to create and analyze dynamic simulations of movement. *IEEE Transactions on Biomedical Engineering*, 54(11), 1940–1950. <https://doi.org/10.1109/TBME.2007.901024> PubMed
36. Todorov, E., Erez, T., & Tassa, Y. (2012). MuJoCo: A physics engine for model-based control. *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033. <https://doi.org/https://doi.org/10.1109/IROS.2012.6386109>
37. Lee, J. D., & See, K. A. (2004, Spring). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. <https://doi.org/10.1518/hfes.46.1.50.30392> PubMed
38. Orozco, J. A., & Artemiadis, P. (2024). Extracting human levels of trust in human–swarm interaction using EEG signals. *IEEE Transactions on Human-Machine Systems*, 54(2), 182–191. <https://doi.org/10.1109/THMS.2024.3356421>
39. Tolmeijer, S., Weiss, A., Hanheide, M., Lindner, F., Powers, T. M., Dixon, C., & Tielman, M. L. (2020). Taxonomy of trust-relevant failures and mitigation strategies. *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 3–12. <https://doi.org/https://doi.org/10.1145/3319502.3374793>
40. DiGiovanna, J., Mahmoudi, B., Fortes, J., Principe, J. C., & Sanchez, J. C. (2009, January). Coadaptive brain-machine interface via reinforcement learning. *IEEE Transactions on Biomedical Engineering*, 56(1), 54–64. <https://doi.org/10.1109/TBME.2008.926699> PubMed
41. Wu, S., Zhang, X., Huang, Y., Chen, S., Shen, X., Principe, J., & Wang, Y. (2023). *Generative neural spike prediction from upstream neural activity via behavioral reinforcement*. bioRxiv. <https://doi.org/https://doi.org/10.1101/2023.07.25.550495>
42. Rodríguez-Fernández, A., Lobo-Prat, J., & Font-Llagunes, J. M. (2021, February 1). Systematic review on wearable lower-limb exoskeletons for gait training in neuromuscular impairments. *Journal of Neuroengineering and Rehabilitation*, 18(1), 22. <https://doi.org/10.1186/s12984-021-00815-5> PubMed
43. Wang, Y., & Principe, J. C. (2021). Reinforcement learning in reproducing kernel hilbert spaces. *IEEE Signal Processing Magazine*, 38(4), 34–45. <https://doi.org/10.1109/MSP.2021.3076309>
44. Philips, G. R., Daly, J. J., & Príncipe, J. C. (2017, July 6). Topographical measures of functional connectivity as biomarkers for post-stroke motor recovery. *Journal of*

Neuroengineering and Rehabilitation, 14(1), 67. <https://doi.org/10.1186/s12984-017-0277-3> PubMed

45. Chambers, V., & Artemiadis, P. (2023, January 4). Using robot-assisted stiffness perturbations to evoke aftereffects useful to post-stroke gait rehabilitation. *Frontiers in Robotics and AI*, 9, 1073746. <https://doi.org/10.3389/frobt.2022.1073746> PubMed
46. Chambers, V., & Artemiadis, P. (2023). A model-based analysis of the effect of repeated unilateral low stiffness perturbations on human gait: toward robot-assisted rehabilitation. *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 12631–12637. <https://doi.org/https://doi.org/10.1109/ICRA48891.2023.10160224>
47. Chambers, V., & Artemiadis, P. (2025). Unilateral compliant surfaces in post-stroke gait retraining: enhancing symmetry and stability. *2025 International Conference On Rehabilitation Robotics (ICORR)*, 267–271. <https://doi.org/https://doi.org/10.1109/ICORR66766.2025.11063138>
48. Hsu, L., Marquez Hernandez, R., McCoy, K., Decker, K., Vemuri, A., Dominick, G., & Heintzelman, M. (2022). Towards development of an automated health coach. In E. Krahmer, K. McCoy, & E. Reiter (Eds.), *Proceedings of the First Workshop on Natural Language Generation in Healthcare* (pp. 27–39). Association for Computational Linguistics. <https://aclanthology.org/2022.nlg4health-1.4/>
49. Koushik, V., & Kane, S. K. (2022). Towards augmented reality coaching for daily routines: Participatory design with individuals with cognitive disabilities and their caregivers. *International Journal of Human-Computer Studies*, 165, 102862. <https://doi.org/10.1016/j.ijhcs.2022.102862>
50. Wilson, S., Roper, A., Marshall, J., Galliers, J., Devane, N., Booth, T., & Woolf, C. (2015). Codesign for people with aphasia through tangible design languages. *CoDesign*, 11(1), 21–34. <https://doi.org/10.1080/15710882.2014.997744>
51. Kane, S. K., Linam-Church, B., Althoff, K., & McCall, D. (2012). What we talk about: Designing a context-aware communication tool for people with aphasia. *Proceedings of the 14th International ACM SIGACCESS Conference on Computers and Accessibility*, 49–56. <https://doi.org/https://doi.org/10.1145/2384916.2384926>
52. Moffatt, K., McGrenere, J., Purves, B., & Klawe, M. (2004). The participatory design of a sound and image enhanced daily planner for people with aphasia. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 407–414. <https://doi.org/https://doi.org/10.1145/985692.985744>
53. *Healthy People 2030*. (n.d.). Social determinants of health. Retrieved from <https://odphp.health.gov/healthypeople/priority-areas/social-determinants-health>
54. Krahm, G. L., Walker, D. K., & Correa-De-Araujo, R. (2015). Persons with disabilities as an unrecognized health disparity population. *American Journal of Public Health*, 105 Suppl 2(Suppl 2), S198-206. <https://doi.org/https://doi.org/10.2105/AJPH.2014.302182>
55. *Accessible Technology*. (2023, August 22). CRA. <https://cra.org/accessible-technology/>
56. Kim, E., & McCoy, K. F. (2018). Multimodal deep learning using images and text for information graphic classification. *Proceedings of the 20th International ACM SIGACCESS*

Conference on Computers and Accessibility, 143–148.
<https://doi.org/https://doi.org/10.1145/3234695.3236357>

57. Tong, X., Mauriello, M. L., Mora-Mendoza, M. A., Prabhu, N., Kim, J. P., & Paredes Castro, P. E. (2023). Just do something: comparing self-proposed and machine-recommended stress interventions among online workers with home sweet office. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–20. <https://doi.org/https://doi.org/10.1145/3544548.3581319>
58. Mauriello, M. L., Tantivasadakarn, N., Mora-Mendoza, M. A., Lincoln, E. T., Hon, G., Nowruzi, P., . . . Paredes, P. E. (2021, September 14). A suite of mobile conversational agents for daily stress management (Popbots): Mixed methods exploratory study. *JMIR Formative Research*, 5(9), e25294. <https://doi.org/10.2196/25294> [PubMed](#)
59. Powers, T. M. (2011). Incremental machine ethics. *IEEE Robotics & Automation Magazine*, 18(1), 51–58. <https://doi.org/10.1109/MRA.2010.940152>

Copyright (c) 2026 Delaware Academy of Medicine and Public Health.

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<https://creativecommons.org/licenses/by-nc-nd/4.0/>) which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.