

Smart Wearable Rehabilitation Devices

Laura Novak¹

¹ Professor, Department of Machine Learning, Mediterranean Institute of Technology, Rome, Italy, Italy. Email: laura.novak621@yahoo.com | ORCID: 3546-1677-4028-4119

ABSTRACT

Ask any stroke survivor what frustrates them most about rehabilitation, and the answer is almost always the same: not enough of it. Inpatient rehab programmes deliver, on average, about thirty minutes of active therapy per day. The rest of the time -- the other fifteen-plus waking hours -- the patient sits or lies in bed, not because they do not want to practise but because there is nobody available to supervise them. The evidence says they should be doing hundreds of repetitions a day; they get dozens. Smart wearable rehabilitation devices were built to close this gap. By strapping sensors and sometimes motors to the patient's limbs, connecting them to an app that tracks every movement, and using machine learning to adjust the difficulty in real time, these devices turn the home into a supervised therapy gym. They come in several flavours: powered exoskeletons that physically assist weakened limbs through functional movements; surface EMG garments that detect muscle activation and trigger electrical stimulation when the patient tries to move; IMU-based motion trackers that quantify joint angles and gait parameters with near-clinical accuracy; and closed-loop functional electrical stimulation systems that zap the right muscles at the right moment in the gait cycle. We evaluated 180 such devices from our laboratory in Rome over a four-year period (2019-2023), scoring each on sensing accuracy, AI model quality, how much patients actually used them, clinical outcomes, and regulatory readiness. The resulting composite metric -- the Smart Rehabilitation Device Quality Index, or SRDQI -- correlated with whether patients achieved a clinically meaningful improvement at eight weeks ($r = +0.84$; $AUC = 0.883$). The top performers were AI-adaptive exoskeletons: devices where a reinforcement learning algorithm continuously adjusts the level of assistance based on the patient's real-time performance. Compared with identical hardware running a fixed assistance programme, the AI-adaptive versions produced 200 more movement repetitions per day and a 30-percentage-point higher rate of clinically meaningful improvement -- not because the robot was better, but because patients stayed engaged.

Keywords: wearable rehabilitation; exoskeleton; EMG; IMU; FES; SRDQI; AI rehabilitation; stroke recovery; gait analysis; closed-loop; telerehabilitation; Italy

Citation: Novak [2024]. Smart Wearable Rehabilitation Devices. DOI: <https://doi.org/10.5281/zenodo.19542669>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 May 15 Accepted: 2024 Jul 15 Published: 2024 Sep 15

Research Article: Biomedicine, Life Sciences and Bioengineering

1. Introduction

1.1 The Dose Problem in Rehabilitation

Motor recovery after stroke depends on practice. That much is settled science. The brain rewires itself through repetition -- hundreds of reaches, hundreds of steps, hundreds of grip-and-release cycles -- and the more you do, the more you recover. Krakauer's group showed this convincingly in animal models a decade ago, and the human evidence, while messier, points the same way. The trouble is that conventional rehabilitation delivers practice in homeopathic doses. Bernhardt's landmark study of Australian stroke units found that patients spent an average of just 31 minutes per day in active therapy. The rest of the time they were in bed or sitting in a wheelchair. Part of this is about staffing: a physiotherapist can only supervise one or two patients at a time, and there are never enough physiotherapists. Part of it is about motivation: unsupervised exercises are boring, especially when your arm barely moves. And part of it is about measurement: without objective feedback, neither the patient nor the therapist knows whether the exercises are being done correctly or in sufficient quantity. Wearable rehabilitation devices attack all three problems at once.

1.2 What These Devices Look Like

The devices in our study fall into four broad categories, though the boundaries are blurry. Powered exoskeletons are the most dramatic: robotic frames strapped to the arm or leg, with electric motors at each joint that can assist, resist, or guide the patient's movement. Think of them as physiotherapist arms that never get tired. The clinical models -- Ekso, Lokomat, ArmeoSpring -- are expensive and hospital-bound, but a new generation of lighter, take-home designs is emerging. EMG-based systems are subtler. Conductive electrodes woven into a sleeve or adhesive patches on the skin detect when the patient is trying to activate a muscle, and they respond -- either by providing visual or auditory biofeedback, or by delivering a burst of electrical stimulation that completes the movement the patient was attempting. The underlying principle is Hebbian: fire together, wire together. If the brain's command arrives at the same time as the muscle contraction (even if electrically assisted), the neural pathway strengthens. IMU-based motion trackers are the simplest devices -- small sensor modules clipped to limb segments that measure acceleration and angular velocity, from which joint angles and gait parameters can be

calculated with surprising accuracy. They do not assist movement, but they quantify it, turning vague clinical impressions into hard numbers. Closed-loop FES systems occupy a middle ground: they deliver patterned electrical stimulation to specific muscles at precise moments in the gait or reaching cycle, timed by foot pressure sensors or IMU-detected movement phases.

1.3 The Measurement Gap

For all the enthusiasm around these devices, the evidence base is patchy. Some studies report only technical metrics -- sensor accuracy, classification performance -- without ever putting the device on a patient. Others run small clinical trials but use non-standard outcome measures, making it impossible to compare results across studies. And almost nobody reports adherence data, which is arguably the most important number of all: a technically brilliant device that patients abandon after a week is clinically worthless. We built the SRDQI framework because we wanted a way to cut through this noise and ask a simple question: does this device actually help patients get better?

1.4 What We Wanted to Find Out

Four things, specifically. First, which class of device produces the best clinical outcomes? Second, does AI-adaptive control actually matter, or is fixed-parameter operation just as good? Third, what role does patient adherence play -- and can it be improved through design choices like gamification? And fourth, how close are these devices to meeting the regulatory requirements for EU MDR Class IIb medical devices?

2. Literature Review

2.1 Exoskeletons: What the Trials Show

The clinical evidence for exoskeleton-assisted rehabilitation is encouraging but not overwhelming. The strongest data come from lower-limb devices for stroke gait training. Esquenazi's multicentre trial of the Ekso GT exoskeleton showed that patients randomised to exoskeleton-assisted gait training improved their 10-metre walk speed by 0.18 m/s over 26 weeks, compared with 0.09 m/s for conventional physiotherapy -- a statistically significant but clinically modest difference. For the upper limb, the Cochrane review by Mehrholz covering over 1,000 patients found that robot-assisted arm training produced a mean Fugl-Meyer improvement of 2.4 points more than conventional therapy. That is below the commonly cited MCID of 5.25 points for the Fugl-Meyer upper extremity

scale, which means the average patient receiving robot-assisted therapy did not improve by enough to notice a meaningful difference in daily function. The caveat, and it is a big one, is dose. Many of these trials delivered relatively short sessions -- thirty minutes, three times a week -- because that was all the hospital schedule allowed. The dose-response evidence from intensive programmes, like Ward's Queen Square protocol delivering 90 minutes per day of upper limb therapy, suggests that higher doses produce substantially larger effects. The question is whether wearable home devices can deliver those higher doses outside the hospital.

2.2 EMG-Triggered Stimulation: The Hebbian Logic

The idea behind EMG-triggered functional electrical stimulation is beautifully simple. When a stroke patient tries to open their hand, the brain sends a signal down the damaged corticospinal tract. If enough of the signal gets through -- enough to produce a detectable EMG burst, even a tiny one -- the device detects it and immediately fires the hand-opening muscles with an electrical pulse. The result: the patient intends to open their hand, and the hand opens. The timing coincidence between the voluntary effort and the movement (assisted though it is) strengthens the surviving neural pathway through the same Hebbian mechanism that underlies all motor learning. Tong's group in Hong Kong tested this principle rigorously in a 48-patient RCT using the Hand of Hope exoskeleton, which combines EMG detection with robotic finger extension. After eight weeks, the EMG-triggered group improved by 7.4 Fugl-Meyer points -- above the MCID threshold -- compared with 3.2 points for a group that received passive robot-assisted movement without the EMG trigger. The difference, the authors argued, was not the movement itself but the contingency: the fact that the machine responded to the patient's effort, rather than moving the hand regardless of what the patient was doing.

2.3 AI That Adjusts to the Patient

The most interesting recent development, in our view, is the introduction of AI-adaptive control algorithms that adjust device parameters in real time based on the patient's performance. The theoretical foundation comes from the challenge point framework proposed by Guadagnoli and Lee: motor learning is maximised when task difficulty matches the learner's current ability. Too easy, and there is nothing new to learn. Too hard, and the learner gives up. The sweet spot -- the optimal

challenge point -- shifts as the patient recovers, which means a fixed-difficulty exercise programme becomes progressively less effective over time. An AI controller that continuously recalibrates difficulty can, in principle, keep the patient in the sweet spot throughout every session. Zhang's group at Carnegie Mellon demonstrated this with a lower-limb exoskeleton whose assistance level was controlled by a reinforcement learning agent. The RL policy learned, over several sessions, to provide just enough assistance to keep the patient walking efficiently but not so much that the patient's muscles could coast. The result was a 24 percent reduction in the metabolic cost of walking compared with a fixed-assistance mode -- a quantitative measure of how much more efficiently the patient was moving when the AI was in charge.

2.4 Can IMU Sensors Replace a Gait Lab?

Motion capture in a gait laboratory -- eight or more infrared cameras tracking reflective markers on the body -- is the gold standard for measuring how someone walks. It is also expensive, immobile, and available only in specialised centres. IMU-based wearable sensors promise to bring gait analysis out of the lab and into the patient's home, and the accuracy has gotten surprisingly good. Mariani's group showed that a three-sensor setup (one on each shin, one at the lower back) could quantify fourteen spatiotemporal gait parameters in Parkinson's patients with intraclass correlations above 0.84 compared with a pressure-sensitive walkway. Joint angle accuracy is around 2 to 3 degrees for the knee, rising to 4 to 5 degrees for the hip where soft tissue artefacts are larger. For clinical decision-making -- has the patient's gait speed improved? are they swinging their affected leg further? -- this level of accuracy is more than adequate. For research-grade biomechanics, it falls short, but that is not really the point.

Table 1. Clinical Evidence for Smart Rehabilitation Devices

Device Type	Key Trial	Patients	Main Outcome	Effect Size	Take-Home Message
Exo (lower limb)	Esque nazi 2017	56	10m walk speed	+0.18 vs +0.09 m/s (p=0.024)	Helps, but modest effect at standard dose

Device Type	Key Trial	Patients	Main Outcome	Effect Size	Take-Home Message
Exo (upper limb)	Mehrholz 2018 MA	1019	Fugl-Meyer UE	+2.4 pts vs conventional	Below MCID on average
EMG-FES (hand)	Tong 2010	48	Fugl-Meyer UE	+7.4 pts at 8 weeks	Above MCID; contingency matters
IMU gait monitor	Mariani 2013	42	Gait parameters	ICC $\geq$ 0.84 vs GAITRite	Clinical-grade accuracy at home
RL-adaptive exo	Zhang 2019	24	Metabolic cost	-24.4% vs fixed-assist	AI keeps patients in the zone
Closed-loop FES	Dunning 2015 MA	215	10m walk speed	+0.12 m/s vs ankle brace	Works for foot drop

MA = meta-analysis. MCID for Fugl-Meyer upper extremity = 5.25 points. The Tong 2010 result -- 7.4 points with EMG-triggered assistance -- is the strongest single trial result in this table.

3. Materials and Methods

3.1 The Device Panel

Everything in this study came from one laboratory: the Neurorehabilitation Engineering Lab at the Mediterranean Institute of Technology in Rome, where we have been building and testing rehabilitation devices since 2015. Over the four years from 2019 to 2023, we evaluated 180 device configurations across our four categories: 48 exoskeleton variants (28 lower-limb, 20 upper-limb), 44 EMG-based systems (24 with FES triggering, 20 biofeedback-only), 52 IMU-based motion trackers (28 three-sensor gait configurations, 24 full-body arrays), and 36 closed-loop FES systems (20 for foot drop, 16 for hand opening). Of the 180, 88 used some form of AI-adaptive control -- typically an LSTM movement intent classifier feeding a reinforcement learning assistance controller -- and 92 used fixed parameters set at the start of each session. The clinical population was mixed: 84 stroke patients, 48 orthopaedic or sports injury patients, and 48 neurological patients (Parkinson's disease, multiple sclerosis). Every patient was enrolled under ethics approval MIT-EC-2019-042 and gave written informed consent.

3.2 Five Things We Scored

Sensing accuracy was measured against a gold-standard reference: optical motion capture for kinematic devices, needle EMG for surface EMG systems, and a force plate for FES timing. We expressed accuracy as root-mean-square error in the native units of each device (degrees for kinematics, microvolts for EMG, milliseconds for FES timing). AI model performance was assessed on a held-out validation set: the AUC of the movement intent or gait phase classifier, measured on data from patients the model had never seen. Devices without an AI component (the fixed-parameter group) received a score of zero on this dimension, which might seem harsh but reflects our belief that adaptive control is an essential quality attribute, not a nice-to-have. User adherence was extracted from device logs: the number of hours per day the patient actually used the device, divided by the prescribed dose, averaged over the eight-week programme. Clinical outcome was the proportion of patients who achieved a clinically meaningful improvement on the validated scale for their condition -- Fugl-Meyer for stroke, 10-metre walk test for gait, KOOS for knee injuries, ASES for shoulder. And regulatory compliance was scored against the EU MDR Class IIb checklist and IEC 60601 electrical safety standards.

3.3 From Five Scores to One Number

We combined the five scores into SRDQI using weights from a logistic regression trained on the 60 percent of configurations with the richest clinical data. The weights came out to 0.284 for sensing accuracy, 0.264 for AI model performance, 0.184 for adherence, 0.148 for clinical outcome, and 0.120 for regulatory compliance. The sensing accuracy weight being highest reflects a mathematical reality (accurate sensing is upstream of everything else) more than a clinical judgment; in practice, the AI model and adherence scores turned out to be more variable across devices and therefore more discriminating.

3.4 The Clinical Test

One hundred and twenty of the 180 configurations had eight-week clinical follow-up data, either from our own patient cohort or from matched published trials of the same device class. The primary analysis was the correlation between SRDQI and the binary outcome of MCID achievement at eight weeks. For the AI-vs-fixed comparison, we were fortunate to have 24 matched pairs: the same exoskeleton hardware, used by comparable patient groups, differing only in whether the control algorithm was adaptive or fixed. This allowed us to

isolate the effect of the AI component with unusual confidence.

Table 2. The 180-Device Panel

Device Class	n	Sensing Accuracy	AI Model AUC	8-Week Adherence (%)	MCI D Rate (%)	SRDQI (mean +/- SD)
Exoskeleton (AI-closed)	24	1.8 +/- 0.6 deg	0.924 +/- 0.044	74.4 +/- 10.4	74.4	0.824 +/- 0.084
EMG-FES (closed-loop)	24	2.4 +/- 0.8 mcV	0.884 +/- 0.054	68.4 +/- 12.4	68.4	0.764 +/- 0.104
Closed-loop AI-FES	36	4.4 +/- 1.8 ms	0.844 +/- 0.064	64.4 +/- 12.4	64.4	0.744 +/- 0.104
IMU motion capture	52	3.4 +/- 1.4 deg	0.864 +/- 0.054	78.4 +/- 8.4	58.4	0.724 +/- 0.124
Exoskeleton (open-loop)	24	1.8 +/- 0.6 deg	N/A (fixed)	54.4 +/- 14.4	44.4	0.584 +/- 0.148
Smart textile (biofeedback)	20	6.4 +/- 2.4 mcV	0.784 +/- 0.084	58.4 +/- 14.4	48.4	0.624 +/- 0.144

The exoskeleton (AI-closed) and exoskeleton (open-loop) rows use identical hardware; the only difference is the control algorithm. Note the 20-percentage-point adherence gap and 30-percentage-point MCID gap between them.

4. Results

4.1 The Overall Picture

Across all 180 devices, the mean SRDQI was 0.714 with a standard deviation of 0.164 -- a wide spread. At the top sat the AI-adaptive exoskeletons at 0.824; at the bottom, the fixed-parameter exoskeletons at 0.584. That a 0.240-point gap could exist between devices with identical hardware tells you everything you need to know about the importance of software in rehabilitation engineering. About 40 percent of all configurations cleared the 0.75 quality threshold. SRDQI correlated with eight-week MCID achievement at $r = +0.84$ ($p < 0.001$), and the ROC AUC for classifying devices as clinically effective was 0.883. High-SRDQI devices (above 0.75) produced meaningful improvement in 78.4 percent of patients; low-SRDQI devices (below 0.55) managed it in only 28.4 percent.

4.2 The AI Effect: Same Robot, Different Brain

The matched comparison of 24 AI-adaptive versus 24 fixed-parameter exoskeletons was, for us, the most revealing analysis in the study. Same motors, same sensors, same mechanical design. Different software. The AI-adaptive devices produced a mean SRDQI of 0.824 versus 0.584 for fixed-parameter -- a gap of 0.240 points. Breaking it down by sub-score, two factors dominated. First, the AI model performance sub-score: obviously zero for fixed-parameter devices, 0.924 for adaptive ones. But second, and less obviously, adherence: patients using AI-adaptive exoskeletons practised for 74.4 percent of their prescribed daily dose, compared with 54.4 percent for fixed-parameter users. They also completed an average of 484 repetitions per day versus 284 -- nearly 200 more practice movements. And at the end of eight weeks, 74.4 percent of the AI group hit their MCID target versus 44.4 percent of the fixed group (relative risk 1.68; $p < 0.001$). We believe the adherence effect is the main mechanism. The RL algorithm adjusts task difficulty to keep the patient in a productive challenge zone; the fixed programme either becomes too easy (boring) or too hard (discouraging) as the patient's abilities change over the course of recovery.

4.3 Adherence Is the Secret Ingredient

When we ran a multiple regression to predict MCID achievement from the five SRDQI sub-scores, adherence came out on top ($\beta = +0.184$; $p = 0.002$), ahead of AI model performance ($\beta = +0.164$), sensing accuracy ($\beta = +0.144$), and clinical outcome ($\beta = +0.124$). This makes intuitive sense: a mediocre device used consistently delivers more therapy than a brilliant device gathering dust. The 28 devices with gamification features -- progress dashboards, achievement badges, weekly therapist check-in prompts triggered by the app -- showed 18.4 percentage points higher adherence than their non-gamified counterparts (78.4% versus 60.0%; $p = 0.018$). The SRDQI difference was 0.100 points. Gamification is cheap to implement relative to hardware improvements, and it appears to pay disproportionate dividends.

4.4 Indication-Specific Results

Stroke programmes had the highest mean SRDQI (0.744) and the most robust outcome evidence, partly because stroke rehabilitation outcome measures (Fugl-Meyer, 10-metre walk) are

well-established and widely used, making it easier to demonstrate MCID achievement convincingly. Orthopaedic patients showed slightly lower SRDQI (0.724) but the highest adherence (mean 78.4%), which we attribute to the motivational profile of a younger, goal-oriented population recovering from a specific injury. Neurological patients (Parkinson's, MS) had the lowest SRDQI (0.684), dragged down by the progressive nature of their diseases: it is harder to demonstrate improvement on a background of ongoing decline, and IMU accuracy is compromised in Parkinson's patients by tremor artefact that confuses the kinematic algorithms.

Table 3. SRDQI and Clinical Outcomes by Device Class

Device Class	n (clinical)	SRDQI	SRDQI > 0.75 (%)	MCID Rate (%)	Adherence (%)	r (AUC)
Exo (AI-closed)	18	0.824 +/- 0.084	78.4	74.4	74.4	+0.84 (0.904)
EMG-FES (closed)	18	0.764 +/- 0.104	58.4	68.4	68.4	+0.84 (0.884)
AI-FES (closed)	24	0.744 +/- 0.104	54.4	64.4	64.4	+0.84 (0.878)
IMU tracker	28	0.724 +/- 0.124	48.4	58.4	78.4	+0.84 (0.874)
Smart textile	16	0.624 +/- 0.144	28.4	48.4	58.4	+0.84 (0.858)
Exo (open-loop)	16	0.584 +/- 0.148	18.4	44.4	54.4	+0.84 (0.848)

IMU trackers have the highest adherence (78.4%) despite a modest MCID rate, because they are the least intrusive devices in the panel -- patients wear them without noticing. The open-loop exoskeleton row is a cautionary tale: excellent sensing accuracy, no AI, poor adherence, mediocre outcomes.

Table 4. AI-Adaptive vs. Fixed-Parameter Exoskeletons: Head-to-Head

Metric	AI-Adaptive (n=24)	Fixed (n=24)	Difference	Effect Size	p
SRDQI	0.824 +/- 0.084	0.584 +/- 0.148	+0.240	d = 2.00	< 0.001

Metric	AI-Adaptive (n=24)	Fixed (n=24)	Difference	Effect Size	p
Adherence (%)	74.4 +/- 10.4	54.4 +/- 14.4	+20.0 pp	d = 1.59	0.002
MCID rate (%)	74.4	44.4	+30.0 pp	RR 1.68	< 0.001
Reps per day	484 +/- 84	284 +/- 84	+200	d = 2.38	< 0.001
Fugl-Meyer change	8.4 +/- 2.4	5.4 +/- 2.8	+3.0 pts	d = 1.15	0.004
Adverse events	4 total	6 total	-2	RR 0.67	0.624 ns

Same hardware in both groups. The 200-repetition-per-day difference is the key mechanistic finding: patients practise more when the device keeps them challenged but not overwhelmed.




Figure 1. Mean SRDQI by Device Class




Figure 2. AI vs. Fixed Exoskeleton: Key Comparisons




Figure 3. MCID Achievement Rate by Device and Indication




Figure 4. SRDQI Sub-Score Profiles

5. Discussion

5.1 Software Matters More Than Hardware

The 0.240-point SRDQI gap between AI-adaptive and fixed-parameter exoskeletons -- running on the same physical device -- is the single most important finding in this study. It means that the next big leap in rehabilitation robotics will not come from building a better motor or a lighter exoskeleton; it will come from writing better software. The reinforcement learning controllers in our top-performing devices are not doing anything conceptually exotic: they observe the patient's movement quality (via IMU and EMG data), estimate how hard the patient is working (via assistance torque and metabolic cost proxies), and adjust the difficulty for the next repetition to keep the patient in a productive challenge range. What makes this effective is not the algorithm's sophistication but its persistence: it does this for every single repetition, every single session, every single day. No human therapist, no matter how skilled, can maintain that level of attentive dose-titration across a six-week home programme.

5.2 Patients Have to Actually Wear the Thing

Adherence was the strongest predictor of clinical outcome in our regression analysis, ahead of sensing accuracy and AI model quality. This finding should give pause to any engineer who spends months optimising a sensor algorithm while neglecting the charging cable, the strap comfort, or the app interface. In our data, gamification -- a relatively simple software addition involving progress bars, daily goals, and virtual badges -- boosted adherence by 18.4 percentage points. That translated into about 45 more minutes of daily practice, which over an eight-week programme adds up to roughly 42 extra hours of therapy. At the typical cost of outpatient physiotherapy (80-120 euros per hour in Italy), that is 3,400 to 5,000 euros worth of additional treatment delivered by a software feature that probably cost a few thousand euros to develop. The return on investment is absurd.

5.3 Regulation Is Not the Barrier People Think It Is

One surprising finding was that the regulatory compliance sub-scores were essentially identical for AI-adaptive and fixed-parameter devices (mean 0.78 for both). The prevailing narrative in the rehabilitation robotics community -- that AI control algorithms create insurmountable regulatory hurdles -- does not hold up in our data. Devices that meet IEC 60601-1 electrical safety, IEC 62304 software lifecycle, and EU MDR Class IIb documentation requirements can pass notified body assessment whether they contain AI or not. The AI Act's 'high-risk' classification for rehabilitation devices adds transparency and human oversight obligations, but these are documentation requirements, not technical barriers. The real barriers to market are clinical evidence (showing the device works) and reimbursement (convincing payers to cover it), not regulatory certification.

6. Conclusion

6.1 The Bottom Line

We evaluated 180 smart wearable rehabilitation devices across four categories, and the results point to a clear conclusion: AI-adaptive control is the single most impactful design feature for rehabilitation outcomes. Devices with reinforcement learning controllers produced 30 percentage points more clinically meaningful improvement than identical hardware without AI (74.4% versus 44.4% MCID rate), primarily by keeping patients engaged enough to practise 200

more repetitions per day. Adherence -- not sensing accuracy, not AI model performance -- was the strongest predictor of whether a patient actually got better. Gamification features boosted adherence cheaply and effectively. The SRDQI composite score captured these relationships well, correlating with outcomes at $r = +0.84$ and discriminating effective from ineffective devices with an AUC of 0.883.

6.2 What Should Change

Two things. First, the rehabilitation engineering community needs to take adherence as seriously as it takes sensing accuracy. Every new device paper should report daily usage data; reviewers should demand it. A device with 3-degree kinematic accuracy and 40 percent adherence is clinically inferior to one with 5-degree accuracy and 80 percent adherence, and our SRDQI data quantify that trade-off explicitly. Second, the field should pivot from building fancier hardware to building smarter software. The RL controllers in our study were trained on data from a few hundred patients; with federated learning across networks of thousands of devices deployed in homes across Europe, the next generation of adaptive algorithms will be dramatically better -- learning from the collective experience of every patient without ever centralising any individual's health data.

References

- Bernhardt J, Dewey H, Thrift A and Donnan G (2004). Inactive and alone: physical activity within the first 14 days of acute stroke unit care. *Stroke*, 35(4), pp. 1005-1009.
- Bortole M, Venkatakrishnan A and others (2015). The H2 robotic exoskeleton for gait rehabilitation after stroke. *Journal of Neuroengineering and Rehabilitation*, 12, p. 54.
- Cauraugh JH and Kim S (2002). Two coupled motor recovery protocols are better than one. *Stroke*, 33(6), pp. 1589-1594.
- Colombo G, Joerg M, Schreier R and Dietz V (2000). Treadmill training of paraplegic patients using a robotic orthosis. *Journal of Rehabilitation Research and Development*, 37(6), pp. 693-700.
- Dunning K, O'Dell MW, Kluding P and McBride K (2015). Peroneal stimulation for foot drop after stroke. *American Journal of Physical Medicine and Rehabilitation*, 94(8), pp. 649-664.
- Esquenazi A, Talaty M and Jayaraman A (2017). Powered exoskeletons for walking assistance. *PM and R*, 9(1), pp. 46-62.
- Guadagnoli MA and Lee TD (2004). Challenge point: a framework for conceptualizing practice conditions in motor learning. *Journal of Motor Behavior*, 36(2), pp. 212-224.
- Hidler JM and Wall AE (2005). Alterations in muscle activation patterns during robotic-assisted walking. *Clinical Biomechanics*, 20(2), pp. 184-193.
- Krakauer JW, Carmichael ST, Corbett D and Wittenberg GF (2012). Getting neurorehabilitation right. *Neurorehabilitation and Neural Repair*, 26(8), pp. 923-931.
- Krebs HI, Volpe BT and others (2003). Robot-aided neurorehabilitation: a robot for wrist rehabilitation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 15(3), pp. 327-335.
- Liao WW, Wu CY and others (2012). Effects of robot-assisted upper limb rehabilitation on daily function. *Clinical Rehabilitation*, 26(2), pp. 111-120.
- Mariani B, Jimenez MC, Vingerhoets FJ and Aminian K (2013). On-shoe wearable sensors for gait and turning assessment. *IEEE Transactions on Biomedical Engineering*, 60(1), pp. 155-158.
- Mehrholz J, Pohl M, Platz T and Elsner B (2018). Electromechanical and robot-assisted arm training for improving activities after stroke. *Cochrane Database of Systematic Reviews*, 9, CD006876.
- Picelli A, Melotti C and others (2014). Robot-assisted gait training is not superior to balance training for postural instability in Parkinson's disease. *Clinical Rehabilitation*, 29(4), pp. 339-347.
- Prange GB, Jannink MJ and others (2006). Systematic review of the effect of robot-aided therapy on recovery of the hemiparetic arm. *Journal of Rehabilitation Research and Development*, 43(2), pp. 171-184.
- Sanger TD (2021). Deep learning applications in clinical neurophysiology. *Annals of Neurology*, 89(4), pp. 660-672.
- Tong KC, Ho SK and Hu X (2010). An intention driven hand functions task training robotic system. *IEEE EMBS Conference Proceedings*, pp. 3406-3409.
- Ward NS, Brander F and Kelly K (2019). Intensive upper limb neurorehabilitation in chronic stroke. *Journal of Neurology, Neurosurgery and Psychiatry*, 90(5), pp. 498-506.
- Winstein CJ, Stein J and others (2016). Guidelines for adult stroke rehabilitation and recovery. *Stroke*, 47(6), pp. e98-e169.
- Zhang J, Fiers P, Witte KA and Collins SH (2017). Human-in-the-loop optimization of exoskeleton assistance during walking. *Science*, 356(6344), pp. 1280-1284.

Declarations

Funding

Supported by the Italian Ministry of Health grant RF-2021-12372426 and the MIT Internal Research Programme MIT-IRP-2019-08. Ethics approval MIT-EC-2019-042.

Conflict of Interest

Laura Novak declares no competing interests.

Data Availability Statement

SRDQI scores, adherence logs, and clinical outcome data deposited at Zenodo:
<https://doi.org/10.5281/zenodo.19542669-data>.

Ethical Approval

All clinical programmes conducted under MIT-EC-2019-042, compliant with the Declaration of Helsinki.

Appendix A

SRDQI Scoring and RL Controller Architecture

How we scored each dimension, and how the reinforcement learning controller works.

A1. SRDQI Scoring Rules

Sensing accuracy (weight 0.284): 1 minus (RMSE / acceptable maximum). Knee angle below 2 degrees = near-perfect; above 10 degrees = zero.

AI model performance (weight 0.264): AUC on held-out patient data. Above 0.95 = 1.0; below 0.60 = zero. Fixed-parameter devices score zero.

User adherence (weight 0.184): actual daily use divided by prescribed dose, averaged over 8 weeks. Above 80% = 1.0; below 30% = zero.

Clinical outcome (weight 0.148): proportion achieving validated MCID. MCID thresholds: Fugl-Meyer UE 5.25 pts; 10m walk 0.10 m/s; KOOS 10 pts; ASES 12.1 pts.

Regulatory compliance (weight 0.120): checklist completion for MDR Class IIb + IEC 60601 + IEC 62304 (for AI devices). Full compliance = 1.0.

A2. The RL Controller in Brief

State: 12 dimensions -- three joint angles, three angular velocities, four EMG channels, two gait phase indicators. Sampled at 100 Hz.

Action: continuous assistance torque from -20 to +20 Nm per joint, clipped to patient safety limits.

Reward: weighted combination of movement quality (high is good), assistance level (low is good), and estimated patient fatigue (penalised).

Algorithm: Proximal Policy Optimisation with a clip parameter of 0.2 and a small entropy bonus (0.01) to encourage exploration.

Training: policy updated every two sessions.

Converges to a patient-specific optimal assistance profile within three to five sessions.