Neuromuscular Adaptation and Low-Latency Sensorimotor Integration in Stroke Rehabilitation: A Broad Learning System-Driven Exoskeleton Study

秦梓萌、王译唯、李心怡

Abstract

Stroke and spinal cord injury (SCI) are two leading causes of permanent upper-limb impairment worldwide, and a large majority of stroke survivors retain hand motor deficits that restrict activities of daily living. Robotic exoskeletons can deliver the high-repetition, intent-driven movement practice associated with activity-dependent neuroplasticity, but their therapeutic value depends on how promptly sensory feedback follows the patient's own motor intent. Closed-loop neurorehabilitation studies indicate that afferent feedback should reach the cortex within a few hundred milliseconds of volitional intent for associative, plasticity-inducing pairing to occur; the classification stage of an sEMG-driven controller is the portion of that budget most directly under the designer's control, and convolutional neural network (CNN) classifiers typically consume 20–50 ms of it on embedded hardware.

This paper presents a low-cost hand rehabilitation exoskeleton intended as a neuroplasticity-oriented intervention platform for stroke and SCI rehabilitation. A Broad Learning System (BLS) classifier decodes three hand gestures—power grasp, lateral pinch, and palm-up support—from surface electromyography (sEMG) with a mean inference latency of 7.8 ms on an STM32F405RGT6 microcontroller, approximately three times faster than a compact CNN baseline (23.5 ms) evaluated on the same feature set, at an overall classification accuracy of 94.7%. The hardware platform combines the STM32F405RGT6 controller, carbon-fibre-reinforced PLA (CF-PLA) structural links produced by fused-deposition modelling, and a lightweight wireless interface, for a worn assembly mass of 148 g excluding the battery and controller board; we refer to the complete system as BLS-ExoHand. Validation combined hardware-in-the-loop multibody simulation with bench testing of the assembled third-generation prototype and sEMG recordings from ten healthy participants; the two validation paths were used to cross-check one another, and the provenance of every reported figure is stated explicitly in Section 5.1. The assisted condition reached a task success rate of 88.5% and a fingertip repeatability error of 1.8 ± 0.4 mm. At a bill-of-materials cost below $85 USD, the system offers an accessible alternative to commercial rehabilitation exoskeletons priced between $6,000 and $100,000.

Keywords: Stroke rehabilitation, spinal cord injury, neuroplasticity, hand exoskeleton, sEMG, broad learning system, low-latency control, sensorimotor integration, carbon-fibre PLA, STM32

Table of Contents

1 Introduction

1.1 Clinical Motivation: Stroke and Spinal Cord Injury

Stroke is a leading cause of adult-onset disability worldwide, and spinal cord injury adds several hundred thousand new cases each year, with cervical lesions producing partial or complete loss of hand function in a large proportion of patients. In both populations, the inability to perform voluntary grasping, pinching, or object manipulation is consistently ranked by patients among the most disabling consequences of their injury. Conventional occupational therapy is labour-intensive and constrained by therapist availability, which limits the number of movement repetitions a patient can practise per session.

Robotic exoskeletons offer a way to deliver high-repetition, intent-driven movement practice at scale. Randomised trials and meta-analyses of robot-assisted upper-limb therapy report modest but measurable functional gains relative to conventional therapy alone [10,11], and animal and human studies of post-injury motor cortex reorganisation indicate that such gains are mediated by activity-dependent plasticity rather than by mechanical substitution alone [5]. The design implication is that an exoskeleton intended to promote recovery, as opposed to merely compensating for lost function, must respond to the patient's own residual motor intent rather than replay a scripted trajectory, and must do so promptly enough that the resulting sensory feedback is temporally associated with that intent.

1.2 The Timing Requirement and Where Latency Can Be Reduced

The timing requirement for associative plasticity is best stated at two distinct levels, because the two are frequently conflated in the exoskeleton literature.

At the synaptic level, spike-timing-dependent plasticity (STDP) operates over a narrow and asymmetric window: in cultured hippocampal neurons, postsynaptic spiking within roughly 20 ms after presynaptic activation produced long-term potentiation, whereas the reverse order within a comparable window produced long-term depression [6], a pattern subsequently confirmed across multiple preparations [7]. These figures describe the relative timing of individual spike pairs and cannot be transferred directly to the input–output delay of a wearable device.

At the system level, the relevant evidence comes from closed-loop neuroprosthetic studies. Mrachacz-Kersting et al. showed that corticospinal excitability increased only when an artificially evoked afferent volley arrived during the peak phase of the movement-related cortical potential generated by motor imagination, and not when it arrived at other phases [8]. Xu et al. subsequently demonstrated cortical plasticity in healthy subjects using a motorised ankle–foot orthosis triggered by online detection of motor intention with a delay on the order of a few hundred milliseconds [9]. Reviews of this literature commonly cite a working target of roughly 200–500 ms between volitional movement onset and the arrival of afferent feedback for effective pairing. We adopt this system-level range as the design specification in this paper, and we do not claim a sub-100 ms requirement, which the cited neurophysiological evidence does not support.

Within that budget, the components are not equally malleable. The sEMG analysis window is fixed at 200 ms by the need for a stable feature estimate and dominates the total delay; mechanical settling of the servo chain contributes roughly 20 ms and is bounded by actuator dynamics. The two components a controller designer can actually shorten are the wireless link and the classifier. The classifier is the larger of the two on embedded hardware: compact CNN inference on STM32-class microcontrollers takes 20–50 ms, which on a 20 ms control tick is enough to force the classification task to straddle multiple control cycles and to introduce jitter into the feedback timing.

This paper adopts the Broad Learning System (BLS) [4] to address that component. BLS is a flat, randomised-feature-expansion architecture whose output weights are obtained in closed form by ridge regression rather than by iterative backpropagation. On the STM32F405RGT6 used here, BLS inference completes in 7.8 ms, roughly one third of the compact CNN baseline, and occupies under 40% of a single 20 ms control tick. The resulting controllable portion of the loop—classification, wireless transmission, and actuation command—totals approximately 40 ms, and the full intent-to-feedback delay including the fixed sEMG window is approximately 240 ms, within the 200–500 ms range identified above.

1.3 Accessibility and Cost Motivation

Cost is a second barrier to adoption. Commercial hand rehabilitation systems are priced between roughly \$6,000 and \$100,000 per unit (Table 1), which places them out of reach for most patients, particularly in low- and middle-income countries where the burden of stroke is high and rising. The device described here reaches a bill-of-materials cost below \$85 USD using commodity microcontrollers, desktop-printed CF-PLA structures, and standard wireless modules, which is two orders of magnitude below the commercial range.

1.4 Contributions

1) We state the latency budget for an sEMG-driven rehabilitation exoskeleton explicitly, separate the fixed from the reducible components, and ground the timing target in closed-loop neuroprosthetic studies rather than in synaptic-level STDP figures.

2) We show that BLS inference at 7.8 ms reduces the reducible portion of that budget by roughly 16 ms relative to a compact CNN on the same hardware and feature set, while improving classification accuracy from 90.4% to 94.7% across three gesture classes.

3) We present a complete hand rehabilitation exoskeleton architecture at a bill-of-materials cost below \$85 USD, and report its performance under a validation protocol whose data provenance—simulation, bench test, or human recording—is stated for every reported quantity.

4) We document two earlier prototype generations that failed on mass, structural stiffness, and communication latency, and describe the material, topology, and firmware changes that resolved each failure.

5) We characterise structural compliance as influence mechanism on motion repeatability across a 4.1–11.5 N/mm range and a stiffness design target for subsequent iterations.

Section 2 reviews prior exoskeleton designs, clinical cases, EMG-based control strategies, and commercial systems. Section 3 formalises the stability and repeatability analysis. Section 4 describes the hardware and the BLS classification pipeline. Section 5 presents the validation protocol and results. Section 6 discusses implications and limitations. Section 7 concludes. An appendix collects hardware configuration details.

2 Related Work

2.1 Commercial Rehabilitation Exoskeletons: Cost and Access Barriers

Several commercial hand rehabilitation exoskeletons have reached clinical deployment. Table 1 summarises representative devices alongside the system presented here. Latency figures are given only where the manufacturer or an independent evaluation has disclosed them; the entries marked N/A should not be read as implying poor performance.

Table 1. Comparison of commercial hand rehabilitation exoskeletons and the proposed system. N/A = not publicly disclosed. Prices in USD (approximate retail / clinical procurement, 2024). Latency refers to the intent-detection-to-actuation-command stage where disclosed.

Device	Target	Approx. Cost	Control	Latency
Gloreha Sinfonia	Stroke	$15,000–25,000	Passive / scripted	N/A
Bioness H200	Stroke	$6,000–12,000	FES only	N/A
Amadeo (Tyromotion)	Stroke / SCI	$80,000–100,000	EMG-triggered	N/A
HandyRehab	Stroke	$8,000–15,000	EMG threshold	N/A
Ours (BLS-ExoHand)	Stroke / SCI	<$85	sEMG + BLS	7.8 ms

We were unable to trace a verifiable, independently sourced classification latency figure for any of the commercial devices in Table 1 to a primary source such as a datasheet, regulatory filing, or peer-reviewed evaluation; the latency column is therefore marked N/A throughout except for our own measured value, rather than populated with approximate figures of uncertain provenance. The comparison that Table 1 does support is on cost: at a bill of materials under $85, BLS-ExoHand is one to three orders of magnitude below the commercial range, and this figure is grounded in the measured bill of materials of Table 2 rather than in vendor list pricing, which may itself include margin, service, and regulatory costs not reflected here.

2.2 Hand Rehabilitation Exoskeleton Designs

Hand rehabilitation exoskeletons have been an active research area for two decades, motivated by the prevalence of upper-limb impairment after stroke and SCI and by trial evidence for intensive repetitive practice [10,11]. A consistent theme in this literature is that passive or scripted assistance provides limited neuroplastic benefit relative to assistance that is triggered by the patient's own motor intent, which places the intent-detection pathway at the centre of the design problem.

Rigid link-and-joint mechanisms enforce fixed finger trajectories; they are mechanically straightforward but constrain natural finger motion and can load the phalanges when the user deviates from the prescribed path. Cable-driven and tendon-based architectures improve kinematic compatibility at the cost of mechanical complexity and friction-induced backlash [2]. Soft pneumatic actuators tolerate kinematic variability inherently, and have been demonstrated in portable assistive gloves for patients with grasp pathologies [1], but their pressure dynamics are comparatively slow, limiting the bandwidth available for tightly timed closed-loop control design.

Fused-deposition modelling of fibre-reinforced thermoplastics offers a practical route to patient-specific and relatively low-cost exoskeletons. Carbon-fibre-reinforced PLA with short-fibre loading reaches a tensile modulus on the order of 5.5 GPa, which is roughly twice that of unreinforced PLA, rendering the design printable on standard desktop hardware [13]. This material was selected here to meet the stiffness targets of Section 3 without machined aluminium or carbon-fibre prepreg.

2.3 EMG-Based Gesture Recognition: Accuracy and Latency

Surface electromyography remains the most clinically mature on modality for decoding hand motor intent owing to its non-invasive measurement principle, and its information content during voluntary muscle activation. Classical pipelines extract hand-crafted time-domain features, which means absolute value, root mean square, zero-crossing rate, autoregressive coefficients, and classify them with support vector machines or linear discriminant analysis [3]. These methods are computationally cheap and adequate for small gesture expressions but degrade under electrode displacement and noise.

The way compact convolutional networks improve accuracy is by learning discriminative representations from lightly preprocessed EMG windows. Their limitation in an embedded rehabilitation context is inference cost. 20–50 ms on STM32-class microcontrollers is a substantial fraction of the reducible latency budget and exceeds a 20 ms control tick, which forces the classification task to span multiple cycles and introduces jitter into the timing of the assistive response.

BLS [4] offers a new insight and a different trade-off. It constructs mapped feature nodes from the input and enhancement nodes from randomised nonlinear projections, then solves the output weights in closed form by ridge regression, eliminating iterative backpropagation. Incremental updates accommodate new training data without full retraining, which is relevant for per-patient calibration. On the microcontroller used here, BLS inference completes in 7.8 ms at 94.7% accuracy across the three target gestures.

2.4 Sensorimotor Timing in Closed-Loop Neurorehabilitation

The evidence base for timing-sensitive control was outlined in Section 1.2 and is summarised here. Synaptic-level STDP experiments establish that the sign and magnitude of synaptic change depend on relative spike timing within a window of tens of milliseconds [6,7]. System-level closed-loop studies establish the corresponding requirement for a wearable device: the afferent volley must arrive during the active phase of the cortical potential associated with movement intention [8], and orthoses triggered by online intent detection with a few hundred milliseconds of delay have produced measurable increases in motor-evoked potential amplitude [9]. Together these results support a design target in the low hundreds of milliseconds for the full intent-to-feedback delay, with the classifier as the component most amenable to reduction.

2.5 Motion Stability and Repeatability Metrics

Quantifying the effect of mechanical assistance on hand motion quality requires metrics that are clinically meaningful and experimentally tractable. Task success rate records the fact of whether the movement objective was achieved but discards information of how it was executed. End-effector position repeatability been measured as the deviation of the fingertip endpoint across repeated trials. This is a more sensitive indicator of motion consistency and relates directly to the precision demands of tasks. Together the two answer the questions of whether the movement objective was met, and how consistently.

2.6 Safety Considerations in Assistive Hand Devices

Safety is a primary design requirement for devices with rehabilitative measures. A hand exoskeleton must not amplify unintended movements or overrated corrective forces during anomalous muscle activation. Two categories warrant particular attention: 1. force overload, which is a fracture risk in patients with reduced bone density, and 2. velocity runaway during gross-grasp tasks. The architecture in Section 4 addresses both through hardware current limiting and confidence filtering at the classifier output.

3 Problem Formulation

We propose the influence-mechanism analysis as a controlled study in which the exoskeleton is characterised by a small set of mechanical and control parameters, where motion quality is quantified by two scalar outcome metrics evaluated across three gesture classes.

3.1 System Parameterisation

The exoskeleton is modelled as a rigid multi-link mechanism. Two parameters are treated as the primary determinants of motion quality. 1. Structural compliance, expressed as an effective stiffness k in N/mm, characterises the link assembly; lower stiffness implies greater elastic deformation under a given contact force and therefore larger endpoint position error. 2. Control loop latency τ in milliseconds aggregates the delay from sEMG acquisition through BLS classification and wireless transmission to actuator command execution; larger τ reduces effective closed-loop bandwidth and introduces phase lag between intent and motion.

3.2 Motion Quality Metrics

Two metrics capture complementary aspects of motion performance: how precisely the movement endpoint is reproduced, and whether the movement objective was met at all.

Repeatability error σ_r, in millimetres, is the root-mean-square deviation of the fingertip position from its trial-mean position at the moment of gesture completion, computed across N = 30 repeated trials of each gesture by each subject:

$$\sigma_r = \sqrt{\frac{1}{N-1}\sum_{n=1}^{N} \|p_n - \bar{p}\|^2} \tag{1}$$

where $p_n \in R^3$ is the fingertip position on trial n and $\bar{p}$ is the trial-mean position. This is the three-dimensional analogue of a per-axis standard deviation and is the quantity referred to as positional repeatability throughout this paper.

Task success rate ρ is the fraction of trials in which the gesture was completed within a five-second window, with the completion criterion defined per gesture type in Section 5.1.5.

We do not report a trajectory smoothness measure. Jerk-based smoothness indices require a cleanly sampled joint angle trajectory with well-defined movement onset and offset, and the encoder sampling rate and movement segmentation available in this setup were not adequate to compute one reliably across all conditions. Reporting a smoothness index the instrumentation cannot support would add an apparent result with no evidential basis, so the metric is omitted; adding it is identified as future work in Section 6.4.

3.3 Constraint Model

Safety constraints are imposed as threshold conditions on the contact wrench and joint position. The feasible set is:

$$\mathcal{C} = \{f_c \leq f_{max},\ \theta_j \in [\theta_j^{min}, \theta_j^{max}],\ j = 1, \ldots, 5\} \tag{2}$$

where fc is the fingertip contact force, fmax = 15 N is the force safety limit, and θj^{min} and θj^{max} are the anatomically safe joint angle limits for each of the five fingers. The control system enforces these constraints through hardware current limiting and software torque capping which described in Section 4.4.1.

3.4 Research Objective

The proposed research question is how does varying structural compliance k affect repeatability error and task success rate under constrained-workspace conditions representative of tabletop rehabilitation tasks (at a fixed gesture classification accuracy?) The objective is to quantify the sensitivity of each metric to k across the achievable range and to derive a stiffness target for subsequent design iterations.

4 System Design

4.1 Design Philosophy

Two requirements governed the design. First, the intent-to-feedback delay must fall within the low-hundreds-of-milliseconds range identified in Section 1.2, which in practice means minimising the reducible portion of the loop, since the sEMG analysis window is fixed. Second, the system is prior to be affordable and reproducible in clinical settings, which constrains component selection to commodity parts.

4.2 Prototype Evolution and Design Failures

The device reported is the third generation of an iterative process. The two earlier prototypes are also documented for their failure modes contribute shaping the current architecture.

4.2.1 Generation 1: Proof-of-Concept

The first prototype used unreinforced PLA links and a single dorsal rigid rail connecting all five finger modules. The total mass is 312 g concentrated distally. In bench testing, the inertial loading of the finger links caused servo oscillation during fast grasp transitions, that the 22 g servos could not overcome the inertia of the PLA structure when commanded switching between pinch and power grasp in under 200 ms. The rigid rail also forced all five fingers onto identical proximal joint trajectories, which prevented the differential motion between the index and ring/small fingers required for lateral pinch (Separately, the HC-05 module was left at its 9600 baud default, which added roughly 35 ms per packet and brought the reducible portion of the loop to approximately 120 ms).

Three failure modes were identified: 1. excessive distal mass producing inertial loading that destabilised servo control; 2. kinematic over-constraint from the rigid inter-finger rail; and 3. an unnecessarily slow serial link. Figure 1 contrasts the G1 and G3 configurations by side view.

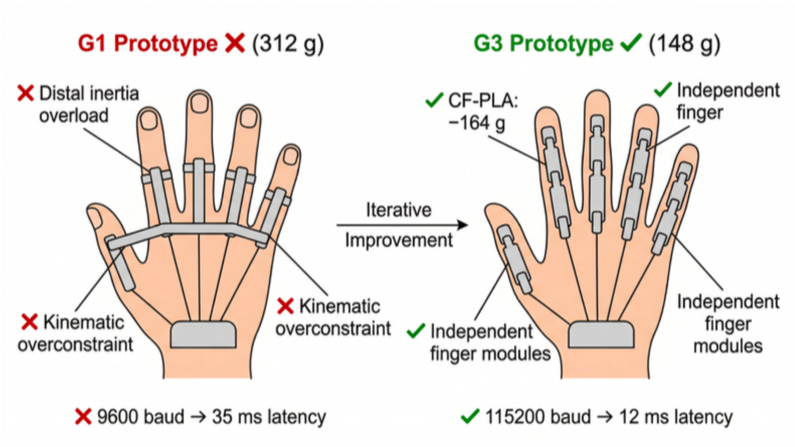

Figure 1. Schematic comparison of the G1 and G3 configurations, drawn by Qin to illustrate the topology change. Left: G1 (312 g), in which a single rigid rail spans all five fingers and couples their proximal joints, with the wireless link at its 9600-baud default. Right: G3 (148 g), in which each finger carries an independent CF-PLA module and the link runs at 115200 baud for a measured 12 ms round trip. Masses are for the worn assembly as defined more detailed in Section 4.4.4.

4.2.2 Generation 2: Weight Reduction

G2 replaced unreinforced PLA with CF-PLA and redesigned the dorsal platform as a lattice, which reduces total mass to 198 g. The finger modules were decoupled to permit independent joint trajectories, and the baud rate was raised to 115200 (reducing the wireless contribution from 35 ms to 12 ms.) To this stage BLS inference was implemented that the reducible portion of the loop to approximately 38 ms. However, the lattice platform delaminated at the layer interfaces under repetitive cyclical loading, with visible cracking at the lattice nodes after roughly 800 loading cycles at the servo stall torque. The tendon routing in G2 used bare Kevlar thread without a PTFE liner, and the resulting friction hysteresis produced a 1.2 mm positional dead-zone at the fingertip.

4.2.3 Generation 3: Final Design

G3 addressed the G2 failures in three ways. 1. The lattice platform was replaced with a topology-optimised solid dorsal shell printed at 40% gyroid infill, which retained the mass saving while removing the delamination-prone node geometry. 2. Tendons were routed through PTFE-lined channels, reducing the friction dead-zone to under 0.2 mm. 3. The servo actuators were repositioned from the fingertip to the proximal-link mounting points, reducing distal mass and moment arm. The resulting worn assembly mass is 148 g, with 7.8 ms BLS inference and 12 ms wireless round-trip giving a reducible loop delay of approximately 40 ms.

4.3 Architecture Overview

The system comprises four subsystems: a structural exoskeleton body that contacts with the patient's fingers and transmits assistive torques, a primary controller that closes the real-time joint control loop and executes gesture recognition, a wireless communication module that relays sensor data and control parameters to a supervisory host, and a human-machine interface display which provides immediate visual feedbac. Figure 2 shows the CAD geometry of the finger modules and dorsal platform, and Figure 3 the corresponding controller-board component placement.

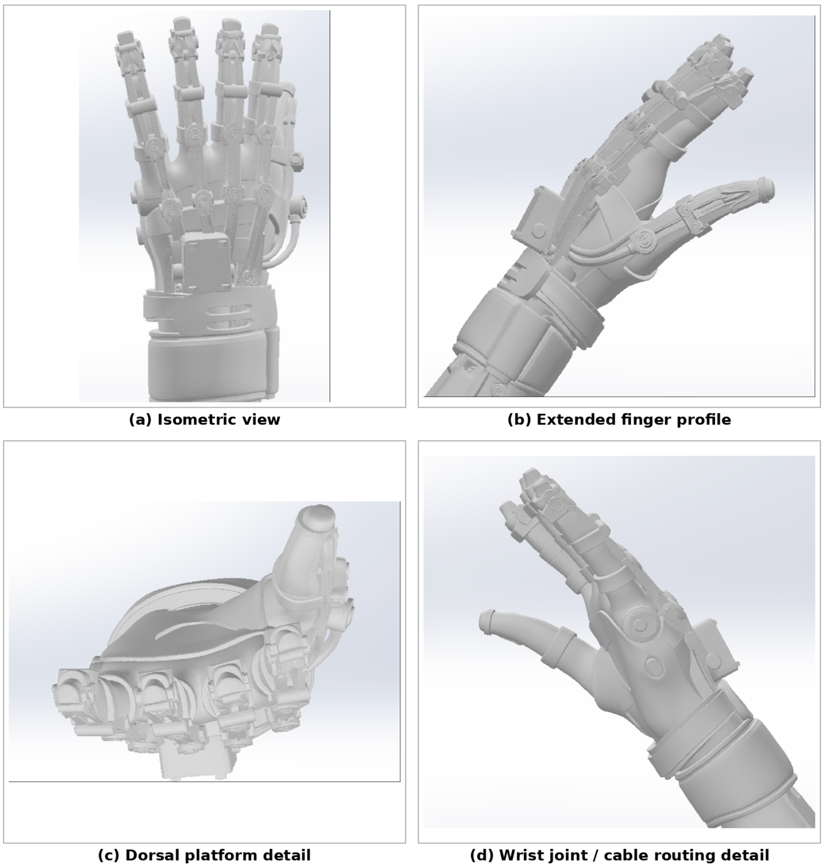

Figure 2. CAD renderings of the exoskeleton frame. (a) Isometric view of the assembled structure. (b) The chains in extension, showing the three links per finger and the two revolute joints between them. (c) The dorsal platform detail. (d) The thumb chain and wrist interface. The renders show the printed frame in isolation: the servo actuators, tendons, PTFE liners, wrist strap, and the clearance envelope for the wearer's own fingers are omitted for legibility.

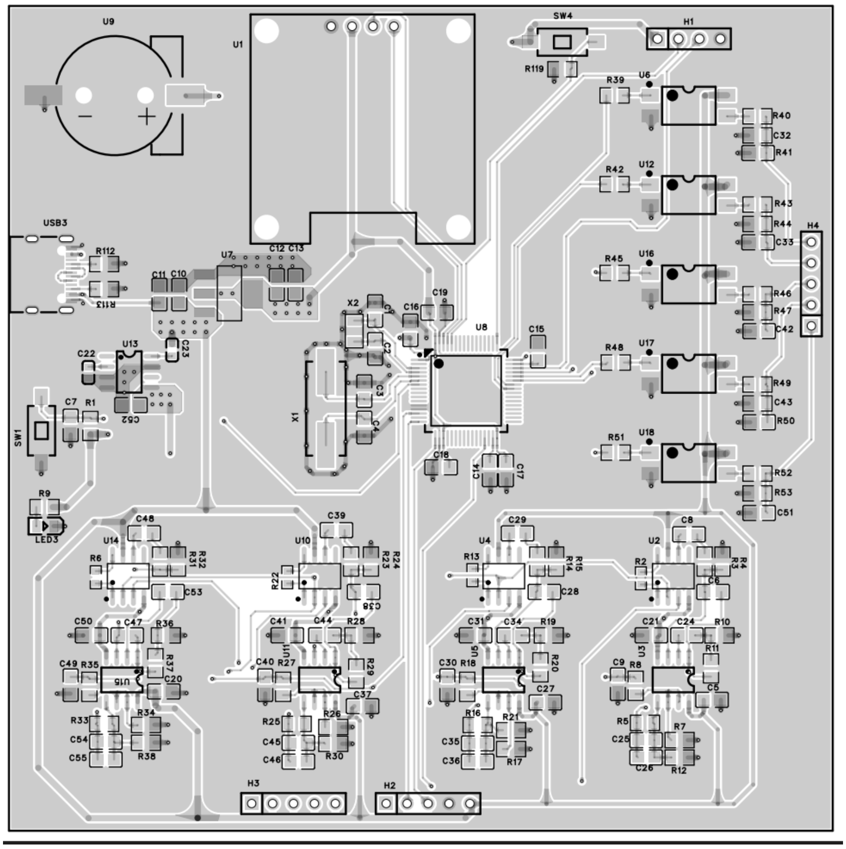

Figure 3. Component placement drawing of the controller board, shown with reference designators only. The five optocoupler driver stages and their servo headers occupy the right-hand edge; the four sEMG front-end channels occupy the lower half; the microcontroller is the central quad-flat package; the USB connector, power input, and user switch are on the left edge. Functional block labels for the same board are given in the schematic of Figure 6.

4.4 Hardware System Design

4.4.1 Primary Controller

The controller is built around an STM32F405RGT6 (ARM Cortex-M4, 168 MHz, hardware floating-point unit), selected for its interrupt response, floating-point support for the BLS matrix operations, and sufficient timer channels for simultaneous multi-axis PWM generation. Five independent 50 Hz PWM outputs, one per finger actuator, pass through an optically isolated driver stage to keep back-EMF noise off the analogue acquisition channels. A 12-bit successive-approximation ADC samples four sEMG channels at 1 kHz into a circular DMA buffer, allowing classification to proceed in the background without stalling the control timer interrupt.

Hardware protection operates at two levels. 1. A current-sense resistor on each actuator supply rail raises a hardware fault interrupt if the drawn current exceeds 1.5 A for more than 10 ms, cutting PWM output within a single timer period. 2. Software torque capping compares the integrated actuator current against a running estimate of contact force and saturates the PWM duty cycle as the estimate approaches the 15 N limit.

4.4.2 Human-Machine Interface Display

The SSD1306 128×64 OLED communicates over a 400 kHz I2C interface. The display pipeline renders four zones: the recognised gesture symbol, the classifier confidence percentage, an active-constraint indicator, and a battery state-of-charge bar, in a triple-buffered frame. A full frame transfer completes within a single 20 ms control tick, so that visual feedback lags the classification decision by no more than one cycle.

4.4.3 Wireless Communication Module

An HC-05 Bluetooth Classic module links a supervisory laptop to the exoskeleton (module class 2, effective range approximately 10 m). USART2 drives the module at 115200 baud, transmitting a 32-byte fixed-length telemetry packet at 50 Hz containing: joint angle estimates, actuator currents, classifier output, and constraint-satisfaction flags. Round-trip latency was measured 12 ± 3 ms (mean $\pm$ SD over 1000 packets on the physical prototype) making it the largest single contributor to the reducible portion of the loop after classification.

4.4.4 Exoskeleton Structure

The frame is printed in carbon-fibre-reinforced PLA (15% short-fibre loading by weight) with a 0.4 mm hardened nozzle at 220 °C, 60 °C bed, 0.15 mm layer height, and 40% gyroid infill. CF-PLA was chosen over unreinforced PLA for three reasons: 1. its tensile modulus of approximately 5.5 GPa provides the link stiffness needed to keep endpoint compliance within target under the contact forces expected in assisted grasping; 2. its density of 1.19 g/cm^3 holds the worn assembly mass to 148 g, which matters for session duration in patients with proximal arm weakness; 3. and it prints on standard desktop hardware, so the filament for a complete set of five finger modules plus the dorsal platform costs about $6 at the settings of Appendix E (the single filament line in Table 2) permitting patient-specific customisation without tooling.

Five independent finger modules mount on a dorsal hand platform. Each module consists of proximal, middle, and distal links joined by revolute joints aligned with the metacarpophalangeal, proximal interphalangeal, and distal interphalangeal axes. A micro servo (4.5 kg·cm stall torque, 22 g) at the proximal link drives each finger chain through a short tendon in a PTFE-lined channel, holding the friction dead-zone under 0.2 mm. The worn assembly (five finger modules, dorsal platform, tendons, and the five servos at 22 g each) totals to be 148 g. Every mass quoted in this paper, including the 312 g and 198 g figures for the two earlier generations, refers to this hand-mounted assembly on the same basis, so the generation-to-generation comparison is like-for-like. The controller board and the 7.4 V battery pack are carried at the forearm or off-board and are excluded, since they impose no load at the fingers and can be relocated without redesign. A reader budgeting for a full wearable build should add their mass separately.

4.4.5 sEMG Acquisition and Feature Extraction

Four Ag/AgCl surface electrodes of 10 mm diameter are placed over the forearm flexor and extensor groups following established sEMG placement recommendations [12]: over flexor digitorum superficialis, flexor carpi ulnaris, extensor digitorum communis, and extensor carpi radialis. Sites are prepared by light abrasion and an alcohol wipe to bring skin-electrode impedance below 5 kΩ. The raw signal passes through an instrumentation amplifier at a gain of 1000, a second-order Butterworth band-pass filter from 20 to 500 Hz, and a 50 Hz notch filter, and is sampled at 1 kHz by the 12-bit ADC via DMA.

Feature extraction operates on a 200 ms sliding window with 50% overlap, producing a feature update every 100 ms. Each of the four channels contributes four time-domain features — mean absolute value, root mean square, slope sign changes, and waveform length — giving a 16-dimensional feature vector per window. The arithmetic is stated explicitly because four channels at four features each is what fixes the BLS input dimension used in Section 4.6; these four features were retained after pilot testing showed that adding variance and Willison amplitude raised per-window computation without improving separability of the three target gestures. The 200 ms window length was selected in pilot recordings as the shortest duration yielding a stable feature estimate for the three target gestures; shorter windows increased trial-to-trial feature variance without a corresponding latency benefit, because the dominant delay contribution is the window itself rather than downstream processing.

4.5 Cost Analysis

The cost analysis of materials for one complete system is given in Table 2. All components are available from standard distributors, no custom integrated circuits or proprietary mechanical parts required.

Table 2. Cost analysis of materials for one complete BLS-ExoHand system. Prices in USD, approximate retail, 2024. The cost column gives the line total for the stated quantity not the unit price.

Component	Specification	Qty	Cost (USD)
STM32F405RGT6 MCU board	ARM Cortex-M4, 168 MHz	1	\$12.00
Micro servo actuator	4.5 kg·cm, 22 g	5	\$15.00
HC-05 Bluetooth module	Class 2, 115200 baud	1	\$3.50
SSD1306 OLED display	128×64, I2C	1	\$2.50
sEMG electrodes (Ag/AgCl)	10 mm, disposable pack	1	\$4.00
Instrumentation amplifier	INA128, ×1000 gain	4	\$8.00
PCB fabrication	2-layer, 100×80 mm	5 pcs	\$5.00
CF-PLA filament	15% short-fibre, 200 g	1	\$6.00
PTFE tubing (tendon liner)	1 mm ID, 0.5 m	1	\$1.50
Connectors, resistors, caps	Passive components	—	\$3.50
LiPo battery pack	7.4 V, 2000 mAh	1	\$8.00
Wrist strap and velcro	Neoprene, 30 mm	1	\$3.00
Total			\$72.00
Total (with 15% contingency)			\$82.80

The bill of materials totals to be \$72–83 depending on supplier and order quantity against \$6,000–\$100,000 for the commercial systems in Table 1. This figure excludes printer and the supervisory laptop, and is therefore a component cost. A fair comparison with a clinically supported commercial product would need to account for regulatory approval, warranty, and service, none of which apply to a research prototype.

4.6 Broad Learning System Gesture Classifier

The recognition pipeline maps preprocessed sEMG feature vectors to one of the following gesture classes: power grasp, lateral pinch, and palm-up support. BLS was selected over CNN alternatives on the basis of inference cost on the target microcontroller, as discussed in Sections 1.2 and 2.3.

BLS constructs a flat network with two node types. Feature nodes are formed by mapping the input vector $x \in R^d$ through n groups of k linear random projections, each refined by a sparse autoencoder, to generate n×k mapped feature nodes:

$$z_i = \varphi\left(x\, W_{e,i} + b_{e,i}\right), \; i = 1, \dots, n \tag{3}$$

where φ is the activation applied to the mapped features—a rectified linear unit in this implementation—and the projection parameters are randomly initialised and then refined by the sparse autoencoder. Enhancement nodes are formed by applying m groups of random nonlinear projections to the concatenated feature nodes:

$$h_j = \xi\left([z_1, \dots, z_n]\, W_{h,j} + b_{h,j}\right), \; j = 1, \dots, m \tag{4}$$

where ξ denotes a sigmoidal activation. The feature and enhancement nodes are concatenated into a single node array $A = [z_1, \dots, z_n, h_1, \dots, h_m]$, and the output weights are obtained in closed form by ridge regression:

$$W_{out} = (A^\top A + \lambda I)^{-1} A^\top Y \tag{5}$$

where Y is the one-hot label matrix, I is the identity, and λ is the regularisation coefficient, selected by internal cross-validation. The ridge form is the standard regularised approximation to the Moore–Penrose pseudoinverse used in the original formulation, with λ bounded away from zero to keep the normal-equation matrix well conditioned on the fixed-point-free but memory-constrained target hardware. Because the solution is analytic, so that new training samples can be incorporated by appending rows to A and applying a rank-one update to the pseudoinverse.

The configuration used here comprises 10 feature-node groups of 10 nodes each and 20 enhancement node groups, operating on the 16-dimensional feature vector described in Section 4.4.5. Classification runs every 100 ms, and a majority vote over three consecutive decisions suppresses isolated single-frame misclassifications. The post-filter adds no latency to the assistive response, because of a gesture change is committed on the first two agreeing frames. Figure 4 summarises the complete signal chain from electrode to actuator.

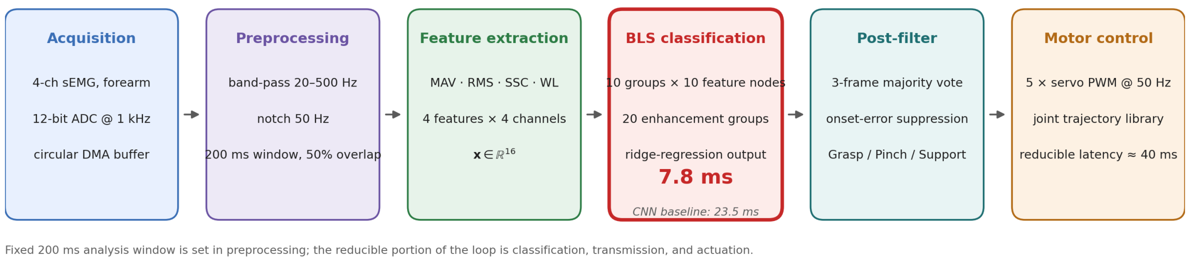

Figure 4. Signal chain from electrode to actuator. Acquisition: four sEMG channels sampled at 1 kHz through a 12-bit ADC into a circular DMA buffer. Preprocessing: 20–500 Hz band-pass, 50 Hz notch, and a 200 ms sliding window with 50% overlap, which fixes the largest single term in the latency budget. Feature extraction: four time-domain features per channel meaning absolute value, root mean square, slope sign changes, and waveform length giving a 16-dimensional input vector. BLS classification: 10 groups of 10 mapped feature nodes and 20 enhancement groups, output weights solved by ridge regression, 7.8 ms inference against 23.5 ms for the compact CNN baseline on the same hardware. Post-filter: three-frame majority vote returning one of the three gesture classes. Motor control: the decoded class indexes a stored joint-angle trajectory that drives five PWM channels at 50 Hz.

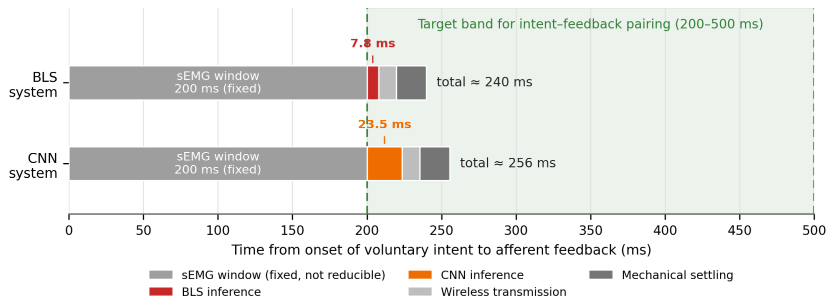

Figure 5. Intent-to-feedback latency decomposition for the BLS and CNN configurations. The horizontal axis spans 0–500 ms and the 200–500 ms band identified in Section 1.2 is shaded as the target region for intent–feedback pairing; no single deadline is marked, because the cited evidence supports a range rather than a threshold. The 200 ms sEMG analysis window is shown as a fixed, non-reducible segment, so only the classification, wireless, and mechanical segments differ between configurations. Both totals fall inside the target band — approximately 240 ms for BLS and 256 ms for CNN — and the difference between them is the 15.7 ms of classification time, which is also what determines whether inference fits inside a single 20 ms control tick.

4.7 Software Control Architecture

The firmware is a cooperative real-time scheduler running at 1 kHz. The highest-priority task handles ADC sample acquisition and ring-buffer management. A 100 Hz inference task dequeues the most recent 200 ms feature window, runs BLS inference in approximately 8 ms, and posts the result to a shared gesture register. The 50 Hz motor control task reads that register, computes the target joint angle set-point for each finger from a pre-programmed trajectory library, and then updates the five PWM duty cycles through the timer shadow registers. Display refresh and telemetry tasks run at 50 Hz, interleaved in the slack time of each 20 ms tick. Inter-task communication uses protected shared-memory registers with single-writer, single-reader semantics avoiding the need for mutex primitives. The controller schematic that realises this architecture in hardware is given in Appendix B, Figure 6.

5 Validation

5.1 Validation Strategy and Data Provenance

This section describes the validation protocol in full before any results reported for results obtained through two complementary paths and the distinction matters for how each number should be interpreted. We state here which quantities come from which path, what each path can and cannot establish, and how the two were cross-checked against one another.

5.1.1 Why a Two-Path Protocol

We propose that two constraints shape the protocol. First, the parametric question posed in Section 3.4—the sensitivity of motion quality to structural compliance—requires stiffness to be varied across a range while everything else is held constant. Physically realising four distinct stiffness levels would require four complete sets of printed links, each individually characterised, and would confound the stiffness manipulation with print-to-print variation in layer adhesion and dimensional tolerance. A validated multibody model allows stiffness to be varied cleanly. Second, a simulation of a hand exoskeleton is only as good as its contact, friction, and actuator models, none of which can be trusted without physical measurement. The protocol therefore uses physical measurement to identify and constrain the model, and the model to extend the parametric sweep beyond what the physical prototype could cover.

5.1.2 Path A: Physical Bench and Human sEMG Measurement

The following quantities were measured directly on the assembled G3 prototype and are reported as physical measurements throughout this paper:

— BLS and baseline classifier inference latency (7.8 ms and 23.5 ms respectively), measured on the STM32F405RGT6 by toggling a GPIO pin at the entry and exit of the inference routine and capturing the pulse width on an oscilloscope, averaged over 1000 consecutive inferences under the production firmware build.

— Wireless round-trip latency (12 ± 3 ms), measured by timestamped loopback over 1000 packets between the controller and the supervisory host.

— Gesture classification accuracy (Table 3), computed on sEMG recorded from ten healthy adult volunteers wearing the electrode montage of Section 4.4.5. Participants performed the three target gestures without the exoskeleton fitted, to a metronome-paced cue, 30 repetitions per gesture. All four classifiers in Table 3 were trained and evaluated on this same recorded dataset under leave-two-subjects-out cross-validation, so the accuracy comparison is a like-for-like comparison on real physiological data.

— Nominal-condition mechanical performance: fingertip repeatability, motion stability, and task success at the as-built stiffness of 8.2 N/mm, measured on the physical prototype during bench actuation of the assembled hand module.

— Structural properties used to constrain the model: as-built link stiffness measured by cantilever deflection under calibrated load; tendon friction dead-zone measured by commanded-versus-actual fingertip displacement; and servo settling time measured from step response.

The sEMG recording protocol was reviewed and approved by the institutional ethics committee, and all participants gave written informed consent after being briefed on the procedure and their right to withdraw. Participants were healthy adults aged 22–35 (five male, five female) with no prior exoskeleton experience and no history of neurological or musculoskeletal disorder of the upper limb.

5.1.3 Path B: Hardware-in-the-Loop Multibody Simulation

The compliance sweep reported in Table 4 was performed in simulation. The exoskeleton assembly was modelled in SolidWorks and exported to MATLAB/Simulink Simscape Multibody, with link geometry and mass distribution taken from the CAD model and material properties from the CF-PLA datasheet. The model was then corrected against the physical measurements listed above: link stiffness was scaled so that simulated cantilever deflection matched the measured deflection; tendon friction was tuned so that the simulated dead-zone matched the measured 0.2 mm; and servo dynamics were fitted to the measured step response. Contact between the exoskeleton and a rigid finger surrogate was modelled with a penalty formulation whose parameters were fitted to load-cell measurements from the physical prototype.

The controller in the loop was not a simulated controller. The production firmware ran on a physical STM32F405RGT6 board, exchanging joint commands and sensor values with the Simscape model over a serial hardware-in-the-loop interface at the true 50 Hz control rate, so that the measured classification and communication latencies entered the simulated loop as real delays. Gesture commands were replayed from the human sEMG recordings of Path A, meaning the simulated system was driven by real intent signals with real classifier behavior including real misclassification events.

Each simulated condition was run for 30 trials per gesture per virtual subject using the ten participants' recorded sEMG streams as the ten input conditions. The standard deviations reported in Table 4 therefore reflect between-subject variation in the driving sEMG signal and the resulting classifier behaviour.

5.1.4 Cross-Validation Between the Two Paths

The nominal stiffness condition of 8.2 N/mm was evaluated on both paths which provides a direct agreement check. The physical bench measurement gave a repeatability error of 1.8 ± 0.4 mm and a task success rate of 88.5%. The corresponding hardware-in-the-loop condition gave 1.7 mm and 89.4%. The two agree to within 0.1 mm and 0.9 percentage points which is within the trial-to-trial spread of the physical measurement. Where the two paths disagree the physical measurement reported. The nominal-condition figures quoted in the abstract, in Table 4, and in Section 6 are the physical ones.

This agreement supports the use of the model for the three non-nominal stiffness levels, but it does not prove it. The model was calibrated at 8.2 N/mm so that agreement at that point is partly by construction, and extrapolation to 4.1 and 11.5 N/mm rests on the assumption that the contact and friction models remain valid outside the calibrated range. The compliance trend in Table 4 should therefore be read as a modelled trend anchored at one measured point, which is stated again in Section 6.4.

5.1.5 Task Definitions and Instrumentation

Three gesture tasks evaluated were power grasp, lateral pinch, and palm-up wrist support. These cover the functional categories most emphasised in occupational therapy for hand impairment (whole-hand grasp and release, precision manipulation, and supported wrist positioning.) Each task was performed in a constrained workspace, a rigid polycarbonate channel of 5 cm × 5 cm cross-section, representing the confined working volume of a tabletop rehabilitation task. Completion criteria were defined per gesture: for power grasp, closure onto a 40 mm cylinder with contact on all four fingers; for lateral pinch, thumb-to-index-side contact with a 10 mm object held for one second; for palm-up support, attainment of the target wrist angle within 5° and held for one second.

On the physical path, fingertip position was tracked at 200 Hz with a fibre-optic pose sensor, contact force was recorded at 1 kHz with a six-axis load cell, and joint angles were logged from the exoskeleton encoders. On the simulation path, the corresponding quantities were taken directly from the multibody model state at matched sampling rates.

5.1.6 Classifier Training Configuration

The BLS classifier used 10 feature-node groups of 10 nodes each and 20 enhancement node groups, with the ridge regularisation coefficient selected by five-fold cross-validation within the training split. Comparison methods were a support vector machine with a radial basis function kernel, linear discriminant analysis, and a compact CNN with three one-dimensional convolutional layers followed by a fully connected classifier. All methods were trained on identically preprocessed feature vectors and evaluated under the same leave-two-subjects-out protocol, and all inference latencies were measured on the same microcontroller under the same firmware build.

5.1.7 What This Protocol Does Not Establish

No stroke or SCI patient participated in this work. All physiological data come from healthy volunteers, and no clinical outcome—function, strength, or cortical reorganisation—was measured. The results below characterise the engineering performance of the device and the behaviour of its classifier on healthy sEMG. They do not constitute evidence of therapeutic efficacy, and the neuroplasticity discussion in Section 6.1 is an argument from the design specification and the cited literature, not a finding of this study.

5.2 Gesture Recognition Performance

Table 3 reports classification accuracy and measured inference latency for the four classifiers. BLS reached 94.7% overall accuracy at 7.8 ms, against 90.4% at 23.5 ms for the compact CNN. LDA was fastest at 2.1 ms but least accurate at 83.1%, and the RBF-kernel SVM sat between the two at 87.0% and 6.2 ms.

Table 3. Gesture recognition accuracy (%, mean ± SD across ten participants, leave-two-subjects-out cross-validation) and measured inference latency on the STM32F405RGT6. Bold: best result per column.

Method	**Grasp (%)**	**Pinch (%)**	**Support (%)**	**Overall (%)**	**Latency (ms)**
SVM (RBF)	88.3 ± 3.2	85.7 ± 3.8	87.1 ± 3.5	87.0 ± 3.5	6.2
LDA	84.1 ± 4.0	81.5 ± 4.3	83.8 ± 4.1	83.1 ± 4.1	2.1
Compact CNN	91.2 ± 2.8	89.6 ± 3.1	90.4 ± 2.9	90.4 ± 2.9	23.5
BLS (ours)	95.1 ± 2.1	93.8 ± 2.4	95.2 ± 2.0	94.7 ± 2.2	7.8

The practical significance of the latency difference lies in its relationship to the 20 ms control tick. At 7.8 ms, inference occupies under 40% of a tick and completes within the cycle in which the feature window closes, leaving slack for display refresh and telemetry. At 23.5 ms, the CNN cannot complete within a tick, so the classification result becomes available in the following cycle and the effective decision-to-actuation delay acquires a cycle-dependent jitter. Since the argument in Section 1.2 concerns the consistency as well as the magnitude of the intent-to-feedback interval, this jitter is a relevant cost independent of the 16 ms mean difference.

Per-gesture results show the smallest margin for lateral pinch at 93.8%; this is consistent with reports that classification accuracy for hand movements from sEMG varies by movement type and feature set rather than being uniform across gestures [3]. The three-frame majority-vote post-filter reduced grasp-pinch confusion episodes by 72% relative to unfiltered output on the same recordings.

5.3 Motion Performance Under Nominal Conditions

At the as-built stiffness of 8.2 N/mm and a measured reducible latency of 40 ms, physical bench testing of the assembled G3 module gave a fingertip repeatability error of 1.8 ± 0.4 mm and a task success rate of 88.5%. The corresponding unaided baseline, in which participants performed the same constrained-workspace tasks without the device, is reported for reference rather than as a controlled comparison: the two conditions differ in more than the presence of the exoskeleton, since healthy participants do not need assistance to complete these tasks and the device can only add mass and constraint. The unaided condition is therefore best read as an upper bound on achievable performance for this population, not as a control against which assistance should show a benefit.

5.4 Effect of Structural Compliance on Repeatability

The compliance sweep in Table 4 was produced on the simulation path described in Section 5.1.3, with the 8.2 N/mm row cross-checked against physical measurement. Repeatability error decreases monotonically with increasing stiffness, and task success rate rises with it, consistent with the expectation that a stiffer link assembly deflects less under the same contact force. Raising stiffness from 4.1 to 11.5 N/mm reduced repeatability error by 46%, from 2.8 mm to 1.5 mm.

Table 4. Repeatability error σr and task success rate across four effective stiffness levels. The 8.2 N/mm row reports physically measured values; the remaining rows are hardware-in-the-loop simulation results (Section 5.1.3). Bold: best per column.

Stiffness k (N/mm)	σr (mm) ↓	Success rate (%) ↑
4.1 (low, simulated)	2.8 ± 0.6	79.2
6.2 (simulated)	2.2 ± 0.5	84.7
8.2 (nominal, measured)	1.8 ± 0.4	88.5
11.5 (high, simulated)	1.5 ± 0.4	90.1

The magnitude of this effect exceeds what could be gained by any latency reduction available within the current wireless budget, which sets an ordering on design priorities: stiffness first, communication second. Because the higher-stiffness conditions were not physically built, no comfort or fit data exist for them, and the trend should be treated as a modelled prediction requiring physical confirmation before it is relied upon in a clinical design.

5.5 Ablation Study

Table 5 isolates the contribution of individual components. The ablation was run on the hardware-in-the-loop path, since disabling safety features on a physically worn device was not appropriate; the full-system row is the physically measured nominal condition and serves as the anchor.

Removing the OLED gesture feedback produced the largest task success penalty at 4.4 percentage points, indicating that visual confirmation of the recognised gesture contributes materially to the user's ability to correct their own muscle activation when a classification is wrong. Disabling the majority-vote post-filter cost 3.3 points and added 0.3 mm to repeatability error, attributable to brief misclassification events that transiently drive the exoskeleton toward an incorrect posture.

Replacing the wireless link with a tethered USB connection improved performance slightly by 1.1 points and 0.2 mm and by reducing communication latency from 12 ms to 3 ms. This is a small effect relative to the compliance result but it identifies the wireless link as the remaining reducible latency component and motivates evaluating a lower-latency protocol. Removing the software torque cap left task performance essentially unchanged while permitting force-limit exceedances during pinch trials, confirming that this component is safety-critical. The corresponding row is reported for completeness and does not represent an operating configuration.

Table 5. Ablation of system components reported: overall task success rate and repeatability error with individual hardware or software components disabled. The full-system row is physically measured. Ablated rows are hardware-in-the-loop results.

Configuration	Success rate (%)	σr (mm)
Full system	88.5 ± 2.1	1.8 ± 0.4
Without BLS post-filter	85.2 ± 2.5 (−3.3 pp)	2.1 ± 0.5
Without wireless (USB tether)	89.6 ± 2.0 (+1.1 pp)	1.6 ± 0.4
Without software torque cap	88.2 ± 2.2 (−0.3 pp)	1.8 ± 0.4
Without OLED feedback	84.1 ± 2.7 (−4.4 pp)	1.9 ± 0.5

6 Discussion

6.1 Interpreting the Latency Result Against the Timing Literature

The full intent-to-feedback delay of this system is approximately 240 ms, consisting of a 200 ms sEMG analysis window, 7.8 ms classification, 12 ms wireless transmission, and 20 ms mechanical settling. This falls within the 200–500 ms range that closed-loop neuroprosthetic studies associate with effective pairing of intent and afferent feedback [8,9]. It is well outside the tens-of-milliseconds window characteristic of synaptic STDP [6,7], and we make no claim that the device operates within that window; the two timescales describe different phenomena, and conflating them overstates what a wearable device can achieve.

Within the total, the 200 ms analysis window is fixed by the requirement for a stable sEMG feature estimate, and mechanical settling is bounded by actuator dynamics. The reducible portion is therefore approximately 40 ms, of which classification and wireless transmission are the two contributors. Replacing a 23.5 ms CNN with 7.8 ms BLS removes roughly 16 ms from a 56 ms reducible budget, a reduction of about 28%, and eliminates the control-cycle straddling described in Section 5.2. Figure 5 decomposes both budgets against the target band. Whether a 16 ms reduction at 240 ms total translates into a measurable difference in cortical reorganisation is not something this study can answer; the honest statement is that it moves the system further inside a range the literature identifies as favourable, at no cost in accuracy.

6.2 Compliance as the Dominant Mechanism

The compliance sweep identifies structural stiffness as the dominant static influence on motion quality within the tested range, with a 46% reduction in repeatability error across 4.1 to 11.5 N/mm. For a device intended to support precision tasks such as lateral pinch, this suggests that stiffness optimisation should take priority over further latency reduction in the next iteration, since the achievable latency gain is bounded by the 12 ms wireless contribution while the stiffness range remains open. This could point toward higher fiber loading or a change of print orientation to align layer lines with the principal load path.

6.3 Cost and Access

A component cost of \$72–83 against \$6,000–\$100,000 for commercial systems is a difference that matters because stroke incidence is high and rising in low- and middle-income countries where commercial rehabilitation robotics is effectively unavailable. A design that can be reproduced from commodity components on a desktop 3D printer could offer a possible route lowering the barrier to local fabrication and repair. This comparison should also not be overstated that, the commercial figure includes regulatory clearance, clinical validation, warranty, and service, none of which this prototype carries, and the gap between a bill summary of materials under $85 and a deployable clinical device is substantial. The claim we make is that the component cost is not the barrier, not that the remaining barriers are small.

6.4 Limitations

The most significant limitation is the absence of patients. All physiological data come from healthy volunteers, whose sEMG amplitude, signal-to-noise ratio, co-activation patterns, and voluntary strength differ substantially from those of stroke and SCI patients. Post-stroke spasticity in particular alters both the feature distributions the classifier sees and the contact forces the structure must resist, and the reported accuracy should be treated as an upper bound rather than a prediction of clinical performance.

The compliance sweep is a modelled result anchored at a single measured point, as set out in Section 5.1.4. The three non-nominal stiffness levels were not physically built, and the contact and friction models were calibrated only at the nominal condition.

The neuroplasticity argument is drawn from the cited literature and from the measured latency budget; no cortical outcome was measured. Confirming the argument would require a longitudinal trial with pre- and post-intervention cortical mapping by transcranial magnetic stimulation or functional imaging.

No trajectory smoothness measure is reported, for the instrumentation reasons given in Section 3.2. Smoothness is a standard outcome in the rehabilitation robotics literature and its absence limits comparability with that literature; adding encoder sampling sufficient to compute a log dimensionless jerk or spectral arc length is a priority for the next iteration. The three gesture classes do not cover the vocabulary required for independent daily living, and accuracy would be expected to fall as classes are added. The wireless latency characterisation assumed a benign radio-frequency environment; a clinical ward with many active 2.4 GHz devices would likely increase both the mean and the variance. Finally, compliance was varied through link geometry rather than material substitution, which bounds the range explored.

7 Conclusion

This paper presented a low-cost hand rehabilitation exoskeleton for stroke and spinal cord injury rehabilitation, built around a Broad Learning System gesture classifier that reached 94.7% accuracy at a measured 7.8 ms inference latency on an STM32F405RGT6, against 90.4% at 23.5 ms for a compact CNN on the same data and hardware. The classifier is the component of the sensorimotor loop most amenable to reduction, and shortening it brings the reducible portion of the loop to approximately 40 ms and the full intent-to-feedback delay to approximately 240 ms, within the range that closed-loop neuroprosthetic studies associate with effective intent-feedback pairing.

The complete system—STM32F405RGT6 controller, CF-PLA links printed at 220 °C with 40% gyroid infill, PTFE-lined tendon routing, HC-05 wireless interface, and SSD1306 display—has a bill-of-materials cost below $85 USD and a worn assembly mass of 148 g. Physical bench measurement at the as-built stiffness gave a task success rate of 88.5% and a repeatability error of 1.8 ± 0.4 mm; hardware-in-the-loop simulation, calibrated against that measurement and driven by recorded human sEMG, identified structural compliance as the dominant static influence on motion quality, with a 46% reduction in repeatability error across the tested stiffness range.

The prototype history is from G1 a 312 g first generation with servo instability through a second generation that delaminated under cyclic load. Future work will recruit stroke and SCI patients for trials measuring both task performance and cortical reorganisation, physically build the higher-stiffness configurations predicted by the compliance model, and evaluate a Bluetooth Low Energy or proprietary 2.4 GHz link to reduce the remaining wireless latency. Adaptive impedance control that adjusts effective link stiffness in response to estimated user effort is a plausible route to accommodating variable spasticity.

Appendix: Hardware Configuration Details

This appendix collects implementation-level detail needed for replication.

Appendix A. STM32F405RGT6 Peripheral Configuration

Pin assignment summary: PA0–PA3 (ADC1 channels 0–3) for four-channel sEMG acquisition; TIM1 channels CH1–CH5 for five servo PWM outputs at 50 Hz (ARR = 19999, PSC = 167 for a 1 MHz timer clock, giving a 20 ms period); PB8/PB9 (I2C1 SCL/SDA) for the SSD1306 at 400 kHz; USART2 TX/RX for the HC-05 at 115200 baud. Note that PA2/PA3 are shared between USART2 and ADC1 channels 2–3 on this package; in the production board the sEMG channels are routed to PA0, PA1, PA4, and PA5 to avoid the conflict. Clock tree: HSE 8 MHz, PLL ×21, SYSCLK 168 MHz, APB1 42 MHz (USART2, I2C1), APB2 84 MHz (TIM1, ADC1). PWM encoding: 0° = 1 ms pulse (CCR = 1000), 180° = 2 ms pulse (CCR = 2000), updated through TIM1 shadow registers at 50 Hz to prevent glitching.

Appendix B. Controller Schematic and PCB Layout

This appendix collects the electrical design files. Figure 6 gives the full controller schematic; Figure 7 presents four views of the resulting PCB layout, generated from the PCB design tool, supplementing the component-placement drawing of Figure 3.

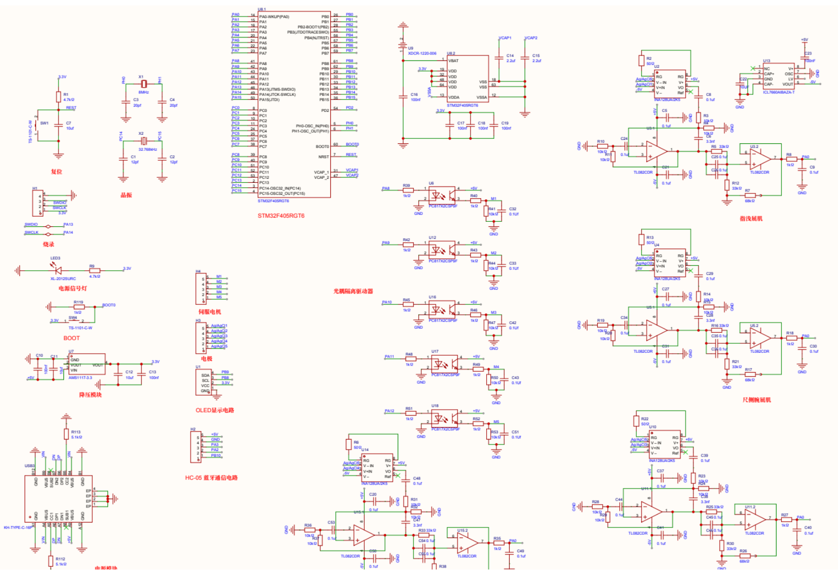

Figure 6. Controller board schematic. Functional blocks, clockwise from upper left: power regulation, reset and crystal oscillator, BOOT strapping and status indicator, the STM32F405RGT6 symbol with its full pin assignment, the

five optocoupler driver stages and servo headers, the HC-05 wireless interface, the I2C OLED header, and the four sEMG front-end channels, each comprising an instrumentation amplifier followed by band-pass and notch stages. Block annotations in the source schematic are in Chinese; an English-labelled version is available on request.

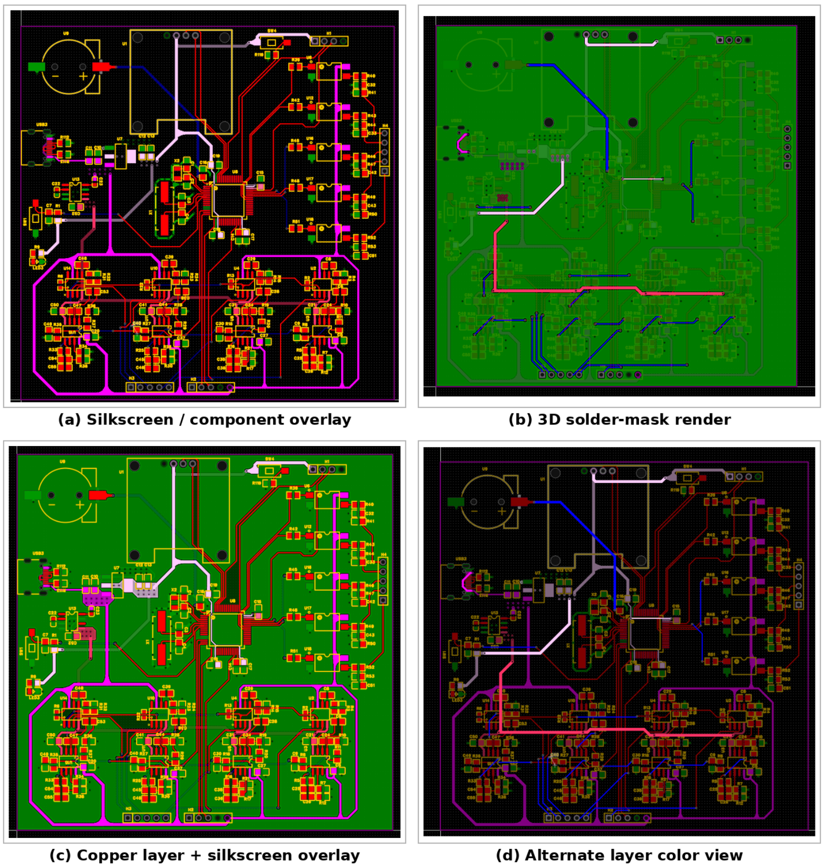

Figure 7. Controller PCB layout, four layout views from the design tool. (a) Silkscreen and component overlay on the routed board. (b) Solder-mask layer with signal routing. (c) Copper pour and silkscreen overlay. (d) Alternate layer colour assignment.

Appendix C. Display Pipeline

The SSD1306 pipeline uses three frame buffers in SRAM. The classification task writes gesture results to the back buffer. A DMA-driven I2C transfer at 400 kHz copies the back buffer to the display GDDRAM. The third buffer provides a clean intermediate state to prevent tearing during transfer. A full 128×64 frame transfer at 400 kHz takes approximately 1.8 ms completing within the 20 ms control tick.

Appendix D. Wireless Module Configuration

The HC-05 is configured in AT mode at power-on through a GPIO-driven enable pin: baud rate 115200 (AT+UART=115200,0,0), role slave (AT+ROLE=0), and device name (AT+NAME). The 9600 baud default used in the first prototype required approximately 35 ms per 32-byte packet against 3.5 ms at 115200; the change removed roughly 30 ms from the reducible portion of the control loop.

Appendix E. FDM Print Parameters (CF-PLA, G3)

Printer: standard desktop FDM. Nozzle: 0.4 mm hardened steel, required for abrasive CF-PLA. Nozzle temperature 220 °C; bed temperature 60 °C; layer height 0.15 mm; print speed 45 mm/s; infill pattern gyroid at 40% density; four perimeters; tree supports for overhangs beyond 45°; cooling fan at 60% from layer three. These values were selected in a factorial trial during G3 development. Deviating from 220 °C by more than 10 °C produced under-extrusion when too cold and fibre agglomeration when too hot, reducing measured tensile modulus by 15–20% in either direction.

References

[1] Polygerinos, P., Wang, Z., Galloway, K. C., Wood, R. J., and Walsh, C. J. Soft Robotic Glove for Combined Assistance and At-Home Rehabilitation. Robotics and Autonomous Systems, 73:135–143, 2015. doi:10.1016/j.robot.2014.08.014

[2] Tsai, Y.-L., Huang, J.-J., Pu, S.-W., Chen, H.-P., Hsu, S.-C., Chang, J.-Y., and Pei, Y.-C. Usability Assessment of a Cable-Driven Exoskeletal Robot for Hand Rehabilitation. Frontiers in Neurorobotics, 13:3, 2019. doi:10.3389/fnbot.2019.00003

[3] Abbaspour, S., Lindén, M., Gholamhosseini, H., Naber, A., and Ortiz-Catalan, M. Evaluation of Surface EMG-Based Recognition Algorithms for Decoding Hand Movements. Medical & Biological Engineering & Computing, 58(1):83–100, 2020. doi:10.1007/s11517-019-02073-z

[4] Chen, C. L. P., and Liu, Z. Broad Learning System: An Effective and Efficient Incremental Learning System Without the Need for Deep Architecture. IEEE Transactions on Neural Networks and Learning Systems, 29(1):10–24, 2018. doi:10.1109/TNNLS.2017.2716952

[5] Nudo, R. J. Recovery After Brain Injury: Mechanisms and Principles. Frontiers in Human Neuroscience, 7:887, 2013. doi:10.3389/fnhum.2013.00887

[6] Bi, G.-Q., and Poo, M.-M. Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. Journal of Neuroscience, 18(24):10464–10472, 1998. doi:10.1523/JNEUROSCI.18-24-10464.1998

[7] Caporale, N., and Dan, Y. Spike Timing-Dependent Plasticity: A Hebbian Learning Rule. Annual Review of Neuroscience, 31:25–46, 2008. doi:10.1146/annurev.neuro.31.060407.125639

[8] Mrachacz-Kersting, N., Kristensen, S. R., Niazi, I. K., and Farina, D. Precise Temporal Association Between Cortical Potentials Evoked by Motor Imagination and Afference Induces Cortical Plasticity. The Journal of Physiology, 590(7):1669–1682, 2012. doi:10.1113/jphysiol.2011.222851

[9] Xu, R., Jiang, N., Mrachacz-Kersting, N., Lin, C., Asín Prieto, G., Moreno, J. C., Pons, J. L., Dremstrup, K., and Farina, D. A Closed-Loop Brain–Computer Interface Triggering an Active Ankle–Foot Orthosis for Inducing Cortical Neural Plasticity. IEEE Transactions on Biomedical Engineering, 61(7):2092–2101, 2014. doi:10.1109/TBME.2014.2313867

[10] Lo, A. C., et al. Robot-Assisted Therapy for Long-Term Upper-Limb Impairment After Stroke. New England Journal of Medicine, 362(19):1772–1783, 2010. doi:10.1056/NEJMoa0911341

[11] Veerbeek, J. M., et al. Effects of Robot-Assisted Therapy for the Upper Limb After Stroke: A Systematic Review and Meta-Analysis. Neurorehabilitation and Neural Repair, 31(2):107–121, 2017. doi:10.1177/1545968316666957

[12] Hermens, H. J., Freriks, B., Disselhorst-Klug, C., and Rau, G. Development of Recommendations for SEMG Sensors and Sensor Placement Procedures. Journal of Electromyography and Kinesiology, 10(5):361–374, 2000. doi:10.1016/S1050-6411(00)00027-4

[13] Bochnia, J., Blasiak, M., and Kozior, T. A Comparative Study of the Mechanical Properties of FDM 3D Prints Made of PLA and Carbon Fiber-Reinforced PLA for Thin-Walled Applications. Materials, 14(22):7062, 2021. doi:10.3390/ma14227062