



Adaptive Orthotic System for Active Hallux Valgus Correction with Real-Time Biomechanical Monitoring

Entisar Younis Abd Al-Jabbar¹, Zaid H. Alsawaff^{1*}, Ahmed Sabeeh Yousif²

¹ Department of Medical Instrumentation Techniques Engineering, Northern Technical University, Technical Engineering College, Mosul 41001, Iraq

² Administrative Technical College, Northern Technical University, Mosul 41001, Iraq

Corresponding Author Email: zaidalsawaff@ntu.edu.iq

Copyright: ©2026 The authors. This article is published by IIETA and is licensed under the CC BY 4.0 license (<http://creativecommons.org/licenses/by/4.0/>).

<https://doi.org/10.18280/jesa.590810>

ABSTRACT

Received: 26 November 2025

Revised: 9 June 2026

Accepted: 18 August 2026

Available online: 31 August 2026

Keywords:

hallux valgus, orthotic devices, biomechanical monitoring, machine learning, wearable electronic devices

Hallux valgus (bunion) is prevalent in adults but conservative treatment continues to be based on passive orthoses, which lack adaptability to patients' biomechanics during wear. In this work, we report a non-invasive orthotic system that continuously tracks hallux alignment, plantar loading, and applied corrective force, allowing real-time adjustments of actuation according to the needs of the user. The system combines an inertial measurement unit (IMU), plantar-pressure sensors and an inline force sensor with edge processing on a Raspberry Pi controller and secure message queuing telemetry transport (MQTT)-based telemetry for remote viewing. Model development was performed on a retrospective dataset of 1,247 sessions from 126 patients, and an independent prospective cohort of 127 patients with mild-to-moderate hallux valgus was used for real-time clinical evaluation. The device applied a corrective force of 0-15 N with 100-ms local control updates and the resolution of alignment was $\pm 0.5^\circ$. In the pilot cohort, active correction demonstrated improved angular alignment and comfort compared to passive bracing, and the magnitude of overloading and under-loading during gait was mitigated. These results demonstrate the potential for patient-specific, non-invasive hallux valgus management, although long-term durability and validation in challenged deformity are still outstanding.

1. INTRODUCTION

Hallux valgus (HV; bunion) is the most common forefoot deformity worldwide. Meta-analysis results indicate that it is estimated to be 19.22% in the adult population with variation according to age, sex and region [1]. The burden follows a sex skewed (female: around 23.74%, male: 11.43%) and age-dependent pattern of increasing among adults and reaching its peak at the age of 60 in adults (%: 22.7 [1]). Longitudinal studies indicate high incidence and progression: about 20% of adults ≥ 50 years will develop HV over 7 years, and about one-third of those with the deformity will get worse [2]. Country-level data are consistent; for instance, in Germany, administrative data show a recorded diagnosis prevalence of ~2%, with about 83% of these being women [3].

Hallux valgus is the lateral deviation of the hallux and medial displacement of the first metatarsal, which results in the formation of a bunion over the first metatarsophalangeal (MTP) joint (Figures 1 and 2). The severity is often assessed by the Manchester scale (none to severe) [4]. The clinical deformity can range from mild angular varus of the hallux to subluxation of the MTP, usually with pronated hallux and associated deformities of the lesser toes.

Passive braces (splints and spacers) provide static support but are not individualized to patient anatomy, tissue response, or changes in deformity, yet they remain widely used for

conservative HV management. Conventional orthoses may cause discomfort, reduce adherence, and offer little or no objective monitoring. Surgery can correct fixed deformity, but it carries perioperative risk, time away from work, and substantial cost; German administrative data also indicate a trend toward decreasing outpatient surgery, emphasizing the need for effective non-invasive options [3].



Figure 1. Manchester scale for grading hallux valgus severity: visual categories (A) none, (B) mild, (C) moderate, (D) severe, employed as a reference standard in clinical evaluations and research reporting [4]



Figure 2. Weight-bearing AP radiograph illustrating key angular measures: (A) hallux valgus angle (HVA: First metatarsal-proximal phalanx); (B) intermetatarsal angle (IMA: First-second metatarsals); (C) distal metatarsal articular angle (DMAA: First-metatarsal shaft-distal articular surface)

A key limitation of currently available braces is the lack of feedback: they cannot evaluate treatment effects during wear, adapt the correction level to individual biomechanics or longer-term patient response, or readily support remote clinical monitoring.

Advances in wearable sensing, secure connectivity, and on-device computation can make HV management more responsive through continuous monitoring and timely intervention [5-8]. Embedded sensors monitor alignment and loading; network connections support secure telemonitoring; and data-driven models can estimate relevant quantities and refine rehabilitation strategies over time [7, 9, 10]. Comparable advances in feedback-enabled medical devices and physiologic closed-loop control illustrate how sensing and actuation can support precise, patient-tailored interventions [11-13].

Based on this, we develop a non-contact orthotic shell that integrates sensing, actuation, connectivity, and local processing within a single monolithic unit. An embedded controller processes inputs from an inertial measurement unit (IMU) for toe alignment, plantar-pressure modules for load distribution, and an inline force sensor for applied correction, and drives a servomechanism that modulates corrective force. Local processing enables 100-ms update intervals, while lightweight messaging supports secure end-to-end remote monitoring and over-the-air software updates. Bench and pilot testing yielded an angular accuracy of $\pm 0.5^\circ$, and the system delivered corrective forces from 0 to 15 N in steady-state operation.

The novelties of this work are: (i) a sensor-to-actuator orthosis that transcends passive support by providing real-time, modifiable correction; (ii) an individualization procedure combining calibrated prediction and incremental adaptation based on observed tolerance and response; (iii) a secure connectivity layer that enables clinician supervision, alerts, and over-the-air updates; (iv) the translation of well-established feedback-control design principles from other medical devices to the domain of orthopaedic rehabilitation; and (v) an evaluation protocol spanning alignment accuracy, force-delivery accuracy, controller stability, and user comfort. In the aggregate, these components outline a strategy for non-invasive, patient-specific control of HV that supplements and potentially reduces dependence on surgical intervention.

2. RELATED WORK

This review positions the present work within three

complementary streams of international research: (i) clinical hallux valgus characterization and grading [1-4], (ii) smart orthoses and wearable sensing for gait and lower-limb rehabilitation [14-16], and (iii) adaptive, data-driven rehabilitation control using supervised learning and reinforcement learning (RL) [17-23]. Earlier clinical studies established how deformity severity, pain, and progression should be measured, whereas more recent engineering work demonstrated that multimodal insoles, ankle-foot orthoses, and connected rehabilitation devices can support continuous biomechanical monitoring across home and clinic settings. This broader perspective is important because the proposed system is not only a hardware prototype; it extends prior sensing and telemonitoring concepts toward localized first-MTP correction.

RL has emerged as a suitable approach for patient-specific control without laborious manual tuning. Molle et al. proposed Q-learning for online adaptation of robot-assisted therapy parameters, selecting stiffness and timing actions from observed kinematics to improve task performance [17]. Pareek et al. [18] extended this concept to assist-as-needed robotic rehabilitation, using a virtual patient to generalize policies and provide smoother, less intrusive assistance than rule-based strategies. Pelosi et al. [19] and co-workers integrated RL with immersive VR by adaptively repositioning targets according to kinematic features and therapist-specified rewards over multi-day training. For hardware involving substantial human-robot interaction, Yang et al. [20] used model-free RL for adaptive admittance control while maintaining smooth transitions between passive and active modes. Collectively, these studies demonstrate the potential of RL for adaptive closed-loop rehabilitation, although their applications were focused on motor training rather than corrective orthoses for structural foot deformities.

Wearable platforms have co-evolved with advances in flexible materials and integration technologies. Kenry et al. [21] reviewed flexible physical sensing platforms for healthcare, emphasizing the trade-offs among sensitivity, mechanical design, stability, and conformability that are critical for skin-mounted applications. Vaghasiya et al. [22] reviewed emerging-material wearable sensors for telehealth and discussed their potential for acquiring physiological indicators with high sensitivity and biocompatibility. Deng et al. [23] reviewed smart wearable systems for health monitoring, highlighting persistent challenges in long-term stability, power management, and system integration for clinical translation. Collectively, these studies support sensor stacks capable of plantar-pressure and alignment monitoring in real-world use, while underscoring the need to address durability, drift, and power consumption.

Three gaps remain to be addressed in this advance. First, most of the published systems are designed to analyze a gait, recognize an activity, or provide a generalized rehabilitation, and not actively correct a forefoot deformity at the first metatarsophalangeal joint. Second, RL-based personalization has been studied mainly for upper-limb or robot-assisted rehabilitation, with scant translation to wearable corrective orthoses. Third, numerous IoT-enabled systems are designed for monitoring and not closed-loop actuation. The present work therefore addresses a smaller, yet clinically motivated challenge: adaptive hallux valgus correction using on-device sensing and local control within the context of remote clinical supervision.

Accordingly, the proposed orthotic platform integrates

multimodal sensing, a supervised force estimator, a safety-constrained reinforcement-learning layer, and a secure clinician-facing IoT pipeline into one application-specific system for first-MTP deformity correction. The aim is not to claim replacement of established clinical pathways, but to demonstrate how adaptive closed-loop correction can complement conservative management and generate more interpretable biomechanical evidence during use.

3. METHODS

For clarity, the Methods section is organized into hardware architecture (Section 3.1), biomechanical modeling (Section 3.2), machine-learning control (Section 3.3), study population and validation strategy (Section 3.4), and IoT deployment (Section 3.5).

3.1 System architecture and hardware components

The orthosis system is developed as a modular three-layer architecture, enabling full-duplex sensing, real-time control, and remote monitoring. (i) Sensing layer: Toe orientation, plantar pressure, and load of correction are sensed. (ii) An embedded processing layer fuses those signals and executes the on-device control policy. (iii) A connectivity layer streams a subset of the metrics for clinician review and enables secure firmware/software updates [24]. The overall system architecture is summarized in Figure 3.




Figure 3. System overview from bottom to top: sensors (MPU6050, FSR-402, HX711), embedded processing on Raspberry Pi 5 (data acquisition, NN inference, DQN reinforcement controller), connectivity (message queuing telemetry transport (MQTT), WebSocket, cloud), and application layer (mobile app, web dashboard). Information flows upward; commands and updates return downstream to close the loop

3.1.1 Hardware configuration

The orthotic prototype was developed using a Raspberry Pi 5 (2.4 GHz ARM Cortex-A76, 8 GB RAM) as the main

controller, accompanied by custom-designed mixed-signal boards for sensor input and actuator output. The system monitors hallux posture, first-ray kinematics, plantar loading, and delivered correction. A 6-axis IMU (MPU6050) tracks the orientation of the hallux and first ray in real time. A single eight-point plantar-pressure array (FSR-402 elements under the first metatarsal and proximal phalanx) evaluates local load distribution associated with valgus correction. Corrective moments about the first metatarsophalangeal axis are generated by three MG996R micro-servos arranged to provide multidirectional actuation. Inline force sensing is implemented with an HX711 amplifier connected to a load cell to measure applied force and close the feedback loop. This configuration enables the device to (i) monitor alignment and localized loading at the first MTP joint and (ii) apply controlled, measurable corrective forces within a compact, wearable form factor suitable for daily use. Representative plantar-pressure distributions across the gait cycle are shown in Figure 4.




Figure 4. Eight-sensor plantar-pressure distribution (0–150 mmHg) visualized as heat maps over six gait phases: Contact, loading response, midstance, terminal stance, pressing, and swing. Each panel shows per-sensor values for a common colour bar

3.2 Biomechanical modeling and mathematical framework

We describe hallux valgus deformity correction as a couple joint-soft tissue system in a first metatarsophalangeal (first-MTP) joint, which depends on kinematics, viscoelasticity of the soft tissues and force-moment considerations. Let the generalized state be:

$$x = [\theta, \theta, \varphi, p_1, p_2, \dots, p_8]^T \quad (1)$$

$$\begin{aligned} \theta_m &= \theta + v\theta, \varphi_m = \varphi + v\varphi, p_m \\ &= [p_1, \dots, p_8]^T + vp \end{aligned} \quad (2)$$

Moments and dynamics.

$$M_{act} = r(\theta, \varphi)u \quad (3)$$

$$M_{pass} = k_1(\theta - \theta_0) + k_3(\theta - \theta_0)^3 + c\dot{\theta} \quad (4)$$

$$I\ddot{\theta} = M_{act} - M_{pass} + d(t) \quad (5)$$

Pressure coupling (discrete CoP shift).

$$\Delta CoP = \sum_{i=1}^8 w_i p_i \quad (6)$$

Discrete-time plant (sample time $T_s = 0.1$ s).

$$X_{k+1} = f(X_k, u_k) + w_k \quad (7)$$

$$y_k = h(X_k) + v_k \quad (8)$$

Safety and comfort constraints.

$$0 \leq u_k \leq 15N, |\theta_k| \leq \theta_{max}, \\ p_{i,k} \leq p_{max}(i = 1, \dots, 8) \quad (9)$$

Control objective (finite-horizon cost).

$$J = \sum_{k=0}^{N-1} [\omega \theta(\theta_k - \theta^*)^2 + \omega F u_k^2 \\ + \omega p \|P_k - P_{safe}\|^2 + \omega \Delta(u_k - u_{k-1})^2] \quad (10)$$

$$u_k = g(X_k) \quad (11)$$

$$r_k = -[\alpha(\theta_k - \theta^*)^2 + \beta u_k^2 \\ + \gamma \text{hotspot}(P_k) + \delta(u_k - u_{k-1})^2] \quad (12)$$

where, θ denotes the hallux valgus alignment angle, θ_0 the painless neutral position, ϕ the hallux-pronation component, p_1 – p_8 the eight plantar-pressure channels, u the commanded corrective force, $r(\cdot)$ the effective moment arm, I the rotational inertia, c the damping term, k_1 and k_3 the linear and cubic stiffness terms, p_{safe} the desired pressure pattern, p_{max} the per-sensor safety cap, and θ^* the target alignment angle. This summary is provided to keep the biomechanical model interpretable for readers from both engineering and clinical backgrounds. The associated Kelvin-Voigt compliance behavior is illustrated in Figure 5.




Figure 5. Biomechanical tissue compliance (Kelvin-Voigt).

(a) Force-displacement response over time (F in N, Δx in mm); (b) hysteresis loop indicating viscoelastic energy loss; (c) session-wise change in effective stiffness K (N/mm) across 20 sessions; (d) mechanical model: linear spring K in parallel with damper C

3.3 Machine learning control framework

An adaptive control system combines supervised learning for state prediction with reinforcement learning for action optimization [24, 25].

In this study, we present a supervised regression model for

short-horizon force prediction based on a retrospective development cohort of 1,247 sessions from 126 patients across mild, moderate, and severe hallux valgus cases. Each session included multiple channels of IMU angles, pressures from eight FSRs, and measurements of applied force at a sampling rate of 100 Hz. To avoid leakage of information, the data was split at the patient rather than window level, with all windows belonging to a single patient being placed in either the training or test subsets. The split ratio stayed at 70%/15%/15% for stratifying according to severity band, sex, and age group to maintain the cohort balance as much as possible. The force-prediction network and representative real-time state signals are shown in Figures 6 and 7, respectively.




Figure 6. Deep neural network architecture for force prediction. Input layer (6 neurons: HVA, $dHVA/dt$, CoP_x , CoP_y , $F_{contact}$, K_{est}) → hidden layers (64 → 32 → 16 neurons with ReLU) → output layer (3 forces). Loss: MSE + L2 regularization ($\lambda = 0.001$)




Figure 7. State-space signals and real-time system behavior: (a) HVA angle; (b) $dHVA/dt$; (c) CoP_x ; (d) CoP_y ; (e) contact force; and (f) hybrid-control force output, all sampled at 100 Hz

The signals were standardized using statistics computed only from the training set and then segmented into overlapping 20-s windows for model development. Sensor data were acquired at 100 Hz, and controller state was updated every 100 ms after temporal aggregation of 10 sensor frames. This distinction decouples sensing frequency from control cadence and will be motivated by the reinforcement-learning timestep

selected in Section 4.3. A shallow temporal encoder with an MLP head was trained using Adam (initial learning rate 1e-3, batch size 32), early stopping, dropout (0.2), weight decay (1e-5), reduce-on-plateau scheduling, and gradient clipping at 1.0.

3.4 Study population, data partitioning, and validation strategy

The development set consisted of retrospective recordings for 126 patients (1,247 sessions; mean age 54.2 ± 12.7 years; 63% female; HVA 12-45°). Severity bands were categorized as mild (12-18°), moderate (18-30°), and severe (>30°), and the same strata were maintained when partitioning the data on a patient basis. Since the purpose of the teaching phase was coverage of biomechanical variability, the inclusion of severe deformities in model building is not to be considered as evidence of clinical efficacy for conservative correction.

Clinical feasibility was subsequently assessed prospectively in an independent cohort of 127 individuals who were not included in model training. This pilot study focused on mild-to-moderate hallux valgus, reflecting the intended conservative-use scenario; individuals with very severe deformity (>40°) were not considered primary candidates for the present protocol. The clinical phase was not randomized and should therefore be interpreted as a feasibility/performance study rather than a definitive comparative trial. Comfort comparisons with passive orthoses were obtained as within-subject control outcomes and were not used for model development.

3.5 IoT framework and mobile application architecture

The solution integrates Raspberry Pi 5-based edge computing, secure cloud services, and a mobile front end to enable real-time monitoring and remote clinical management. A three-tier architecture is used: a device tier acquires and fuses IMU-based alignment, eight-point plantar-pressure, and inline-force signals; a connectivity tier provides low-latency messaging; and an application tier delivers dashboards and patient interfaces for daily use [26].

MQTT, a lightweight publish-subscribe messaging protocol standardized by OASIS [27], is used for device communication. The device publishes biomechanical telemetry (hallux valgus angle, pressure, and force) on sensor topics, listens for high-level model-management commands on a control topic, and reports status on dedicated telemetry channels. Payloads are transmitted as timestamped JSON objects containing a CRC32 checksum. The feedback loop runs locally on the Raspberry Pi rather than through the cloud; therefore, temporary network delay or packet loss can affect dashboard freshness and upload latency but cannot interrupt the 100-ms actuation cycle. During communication loss, the device switches to buffer-and-forward mode and maintains conservative local regulation until the link is restored [28].

A native iOS/Android application displays real-time alignment, a pressure-distribution heat map, and correction force. It collects patient input (comfort score 0-10, notes, and pain scores) and provides a clinician dashboard for remote review, adherence monitoring, and force suggestions [26]. The application connects to the device through a secure WebSocket. For the stated parameters (100-Hz sampling, 24 bytes per data point, and a gzip compression ratio of approximately 0.4), the calculated upstream data rate is 7.7 kb/s; an effective stream of approximately 96 kb/s would

correspond to about 300 bytes per data point at the same sampling rate and compression ratio.

Models can be quantized from 32-bit floating point to 8-bit integer, reducing memory usage by approximately 75% while retaining about 94% of full-precision accuracy. The average inference time was 45 ± 8 ms on the Raspberry Pi 5, below the 100-ms control period [29]. The cloud tier uses AWS IoT Core for secure device connectivity and MQTT brokering, DynamoDB for time-series storage of patient sessions and model performance, and Lambda functions for background tasks such as long-horizon adaptation and monthly compliance reports [28]. Edge-cloud synchronization is adaptive: the base synchronization interval increases as round-trip latency rises to conserve bandwidth and decreases as network conditions improve to maintain data freshness [28].

System validation was performed at the bench level, in simulation, and in a clinical pilot. Bench tests used calibrated standards to validate sensing and actuation. IMU-based alignment was evaluated against reference angles (0° , $\pm 5^\circ$, $\pm 10^\circ$, and $\pm 15^\circ$), yielding a root-mean-square angular error of $\pm 0.5^\circ$. Load-cell force measurements were linear over 0-15 N ($R^2 = 0.9987$), and the eight-sensor plantar array showed $\leq 5\%$ inter-sensor variability under uniform loading. Simulation studies then exercised the full control stack with randomized tissue and noise parameters to assess robustness, followed by a brief clinical pilot to assess supervised-use data quality.

4. RESULTS AND DISCUSSION

4.1 Final system design and implementation

The IoT-connected orthotic solution was realized at the hardware, embedded-control, and cloud-service levels. The final iteration uses three MG996R servomotors controlled by a Raspberry Pi 5 (2.4 GHz ARM Cortex-A76, 8 GB RAM). Biomechanical feedback is acquired at 100 Hz using an MPU6050 IMU (± 3 g accelerometer, $\pm 250^\circ/\text{s}$ gyroscope) with $\pm 0.5^\circ$ angular accuracy, an eight-point FSR matrix for plantar pressure (0-150 mmHg range, $< 5\%$ inter-sensor variation), and an HX711-based load cell (0-20 kg capacity, ± 0.01 kg resolution). Power consumption during active correction is 2.8 W, allowing slightly more than 8 h of continuous operation with a 3000 mAh Li-Po battery. The complete orthotic frame, including the battery, sensors, and actuators, weighs 340 g, and the enclosure measures 280 mm $\times$ 95 mm $\times$ 65 mm.

Edge deployment with TensorFlow Lite eliminated dependence on the cloud and enabled real-time inference. After 32-bit-to-8-bit quantization, model size decreased from 4.2 MB to 1.1 MB (73.8% reduction), with approximately 94% accuracy parity relative to full precision [29]. Mean $\pm$ SD latency was 8.5 ± 1.2 ms ($n = 10,000$) for the regression model and 12.3 ± 1.5 ms for the DQN policy, well below the 100-ms control budget. A real-time scheduler maintained deterministic execution with < 2 ms cycle-time jitter during 72 h of continuous operation; RAM usage was 68% (5.4 GB/8 GB).

4.2 Deep learning model performance on > 1,000 trained cases

In this section, the estimator was trained on the retrospective cohort consisting of 1,247 sessions across 126 patients (mean age 54.2 ± 12.7 years; 63% female; HVA 12-45°) at about

2.494 million time points. These training data were divided at the patient level into training, validation, and test sets (70%/15%/15%) in order to prevent any overlap in the development and evaluation windows [30]. The independent prospective pilot cohort described in Sections 4.4 and 4.5 was excluded from model training.

The training loss decreased rapidly during the early epochs and then approached a plateau. It fell from 1.523 at epoch 1 to 0.073 at epoch 128. The learning curve was well approximated by an exponential-decay model, $L(e) \approx 1.58 \exp(-e/31.2) +$

0.08.

The validation loss followed a similar trend, decreasing from 1.678 to 0.095, with an approximately 30.1% relative development–validation gap at the end of training. The learning-rate schedule was $\alpha(e) = 0.001 \exp(-e/30)$, reaching 5.2×10^{-6} at the final epoch. The average momentum norm was 0.087 ± 0.015 , and the batch-loss coefficient of variation was 8.3%, indicating stable optimization [31, 32]. Table 1 summarizes the cross-validation, residual, and subgroup diagnostics.

Table 1. Summary of model performance and diagnostics: five-fold cross-validation MAEs; residual diagnostics (Shapiro-Wilk and Durbin-Watson); and stratified results by deformity severity (HVA bands), sex, and age

Domain	Subgroup	Sample Size (n)	MAE (N)	R ²	Other Metrics/Notes
Cross-validation (5-fold)	Fold 1	—	0.189 ± 0.041	—	—
	Fold 2	—	0.191 ± 0.047	—	—
	Fold 3	—	0.185 ± 0.043	—	—
	Fold 4	—	0.193 ± 0.049	—	—
	Fold 5	—	0.188 ± 0.045	—	—
Residual diagnostics	Shapiro–Wilk (normality)	373,1	—	—	(p = 0.087)
	Durbin–Watson	—	—	—	1.94
Severity (HVA)	Mild (12–18°)	—	0.156 ± 0.032	0.9741	—
	Moderate (18–30°)	—	0.198 ± 0.051	0.9587	—
	Severe (>30°)	—	0.241 ± 0.068	0.9423	—
Sex	Male	462,7	—	0.9589	—
	Female	743,1	—	0.9651	(p = 0.387) (m vs. f)
Age	Age bands	—	—	0.9521–0.9634	R ² range across age strata

4.3 Reinforcement learning convergence and policy optimization

The DQN was trained with Double DQN, experience replay, target updates every 500 steps, and a soft-update coefficient of 0.001 [33]. Sensor streams were obtained at 100 Hz, but the RL state-action was updated every 100 ms by aggregating 10 sensor frames; thus, an episode of 5,000 control steps is equal to 500 s (approximately 8.3 min). For this definition, the Q-values increased from 15.2 to 94.7, the cumulative reward increased from −45.8 to 287.3, and the TD error was decreased from 38.5 to 1.12 (97.1% reduction). Figure 8 is clinically important because the concurrent increase in reward and decrease in TD error means that the selected terms for reward indeed succeeded in mitigating oscillatory behaviour rather than only motivating the controller to exert higher forces. The RL convergence behavior is summarized in Figure 8.

The exploration schedule $\epsilon(e) = 0.3 \times 0.9995^e$ decreased to $\epsilon = 0.1$ at episode 462 and to $\epsilon = 0.05$ at episode 1020. Across episodes 1–462, random actions accounted for 43.2% of all control steps, whereas exploitation reached 98.1% during episodes 500–1500. In the final 100 episodes, the action proportions for {−2, −1, 0, +1, +2} N were 8.3%, 24.1%, 35.2%, 22.4%, and 10.0%, respectively, indicating a preference for moderate corrections and reduced oscillatory control [33].

Learned policies were conditioned on the estimated stiffness (K) in the range of 2.1–7.8 N/mm. Selected force averaged 4.2 ± 1.1 N for soft tissue ($K < 3.5$ N/mm), 6.1 ± 1.3 N for medium tissue ($3.5 \leq K < 5.5$ N/mm) and 7.8 ± 1.4 N for hard tissue ($K \geq 5.5$ N/mm), which was as expected from biomechanical considerations [34]. Policies trained on 80% of patients transferred with $89.2 \pm 4.7\%$ reward parity to held-out patients, and to adapt to a new patient took 47 ± 18 episodes to get to 95% of baseline performance. Taken together with the action histogram in Figure 8, these results indicate a

preference for intermediate-sized corrective steps and explain why no oscillatory servo behaviour could be seen in deployment.




Figure 8. Convergence analysis of the RL agent: (a) evolution of estimated Q-values; (b) cumulative reward across 1000 episodes; (c) ϵ -greedy exploration schedule, $\epsilon(e) = 0.3 \times 0.9995^e$; and (d) decline of the Bellman TD error

4.4 Real-time closed-loop control performance

The 60% NN/40% RL hybrid controller was subsequently tested in the independent prospective cohort of 127 subjects (mean duration, 4.3 h per subject; 546 h overall). Force-tracking RMSE was 0.237 ± 0.089 N; 95.8% of hallux-angle corrections remained within $\pm 2^\circ$; no safety-bound violations or servo oscillations were detected over more than 219 million control steps; end-to-end response time was 32.4 ± 4.7 ms (95% CI: 28–42 ms); and 99.8% of control cycles were completed within the 100 ± 2 ms execution window. Training diagnostics and servo-actuator performance are summarized in

Figures 9 and 10, respectively.




Figure 9. Neural network training metrics over 150 epochs: (a) training and validation loss convergence (MSE = 0.08); (b) MAE and RMSE; (c) adaptive learning-rate decay (0.001 → 10⁻⁶); and (d) test-set R² reaching 0.95




Figure 10. Servo actuator performance: (a) step response (0 → 10 N; rise time, 1.5 s); (b) frequency response (Bode plot; -3 dB at 50 Hz); (c) multilevel control tracking (RMS error < 0.3 N); and (d) real-time tracking error (mean, 0.235 N)




Figure 11. Integrated system architecture
Note: Top: Sensor layer 8-point FSR array, MPU6050 IMU, and HX711 load cell sampling at 100 Hz over I²C/SPI. Middle: Processor layer: Raspberry Pi 5 (8 GB, 2.4 GHz Cortex-A76) running a 100 Hz control loop with TensorFlow Lite inference. Bottom: Actuation layer three MG996R servos (medial for hallux correction, central for force distribution, lateral for toe support) and a 3000 mAh Li-Po battery (~8 h operation).



Figure 12. Physical device specifications and dimensions. Left: Dimensional profile (280 mm × 95 mm × 65 mm) with color-coded annotations (red = length, blue = depth, green = profile)

Note: Right: Specifications mass 340 g; aluminum-alloy frame; IP54; operating temperature 0-45 °C; active power 2.8 W; 3000 mAh Li-Po (~8 h use; ~2.5 h charge); dual-band Wi-Fi 802.11ac and Bluetooth 5.0 (~50 m indoor/~100 m outdoor); USB-C charging and UART debug; 100 Hz sensor sampling over I²C/SPI; response latency 32.4 ± 4.7 ms.

The mean corrective force was 6.2 ± 1.8 N (range, 2.1-14.8 N). Force profiles showed an adaptation phase during the first 5 min (7.1 ± 2.1 N), a stable phase from 5 to 60 min (5.9 ± 1.4 N), and a later phase after 60 min (5.2 ± 1.1 N), consistent with progressive viscoelastic accommodation [35]. Step-to-step fluctuations were small (≤0.5 N in 94.2% of steps and ≤1.0 N in 99.7%), consistent with stable controller behavior and limited pressure hotspots. Mean comfort was 7.8 ± 1.2 with the active orthosis versus 5.1 ± 1.8 with the passive orthosis in the within-subject baseline condition (paired t = 18.4, p < 0.001; n = 127). This direction of improvement is consistent with systematic-review evidence that hallux-valgus orthoses can reduce pain and improve alignment [36]. The integrated hardware architecture and physical device specifications are presented in Figures 11 and 12, respectively.

4.5 System validation and clinical outcomes




Figure 13. Hierarchical real-time control stack with 100 Hz sensing and 100-ms control updates

Note: Inputs (FSR/IMU/load cell/battery) are normalized to [-1, 1] and processed in parallel by a supervised NN (6 inputs; 64 → 32 → 16; 60% weight) and a DQN policy (40% weight). The hybrid command $F_{cmd} = 0.6F_{NN} + 0.4F_{RL}$ is constrained to 0-15 N and applied to three MG996R servos. Stage latencies are 0-2 ms for sensing, 8.5 ms for NN inference, 12.3 ms for RL inference, 2 ms for fusion, and 3 ms for actuation; aggregate monitoring showed zero safety-limit breaches over 219 million control steps.

In the independent pilot cohort (127 patients; 546 h of use), hallux-valgus angle correction increased over time: $2.1 \pm 0.8^\circ$ at 0.5 h ($78\% \geq 1.5^\circ$), $4.3 \pm 1.2^\circ$ at 1 h ($89\% \geq 3.0^\circ$), $6.8 \pm 1.5^\circ$ at 2 h ($94\% \geq 5.0^\circ$), and $8.4 \pm 2.1^\circ$ at 4 h ($91\% \geq 6.0^\circ$). The maximum observed correction was 12.7° in a moderate case during a 4-h session. The time course showed a rapid early phase followed by smaller additional angular corrections after 1 h, consistent with a saturating viscoelastic adaptation process rather than a linear increase with time. Residual correction was $6.2 \pm 2.1^\circ$ (72% of peak) at 30 min after intervention. Pain scores (VAS, 0–100 mm) decreased from 65.3 ± 18.4 mm at baseline to 38.1 ± 14.7 mm after 4 h (41% reduction; $t = 24.8$, $p < 0.001$). Mild transient erythema occurred in 3 patients (2.4%) and resolved within 2 h without intervention. These findings are consistent with the broader literature on wearable sensing systems for rehabilitation and remote monitoring [37]. The hierarchical control stack and aggregate hybrid-controller performance are summarized in Figures 13 and 14, respectively.




Figure 14. Hybrid controller performance (60% NN, 40% RL) across 127 patients and 546 hours (219 M control steps)
Note: Accuracy/precision: force-tracking RMSE 0.237 N, angular precision 95.8%, 0 servo oscillations. Safety/reliability: 0 safety violations, tracking stability 100%. Timing: mean response 32.4 ± 4.7 ms (28–42 ms); control-cycle compliance 99.8%, target window hit 100%, completion success 99.8%. Compliance: angle within $\pm 2^\circ$ 95.8%, force-tracking accuracy 97.6%. Architecture: 60% supervised estimator, 40% adaptive path. RMSE distribution (N): mean 0.237, +1 SD 0.326, -1 SD 0.148. Overall status: Real-time control with stable tracking and no safety events.

To our knowledge, this is one of the first studies to integrate continuous biomechanical sensing, on-device inference, and closed-loop corrective actuation for hallux valgus management in a wearable orthosis. The primary technical contribution is the combination of accurate force estimation with a clear separation between retrospective model development and prospective pilot validation. This separation supports a more informative assessment of generalizability than a same-cohort analysis. Clinically, the progressive improvement shown in Figure 15 and the residual correction after device removal are consistent with short-term tissue accommodation rather than transient over-tightening alone. The adaptive controller also exhibited a practically meaningful convergence pattern. The reward function penalized angle error, applied force, pressure hotspots, and abrupt action changes; consequently, the learned policy favored gradual corrective adjustments and avoided large step changes that could cause discomfort. Figures 8 and 14 are therefore best interpreted together: convergence of reward and TD error during training translated into stable force delivery,

no detected servo oscillations, and consistent execution timing in the clinical pilot. A large-scale systematic hyperparameter-sensitivity analysis was not conducted and remains an important direction for future work.




Figure 15. Clinical outcomes across 127 patients and 546 h of use
Note: Mean HVA correction reached 8.4° at 4 h (max 12.7°); success rates were 78% ($\geq 1.5^\circ$ at 0.5 h), 89% ($\geq 3.0^\circ$ at 1 h), 94% ($\geq 5.0^\circ$ at 2 h), and 91% ($\geq 6.0^\circ$ at 4 h). Post-treatment retention at 30 min averaged 72% of peak ($\sim 6.2^\circ$). Pain (VAS) decreased 41% ($65.3 \rightarrow 38.1$ mm). Adverse events: 2.4% mild, transient erythema; no serious events. A fitted model indicates an asymptote of $\sim 12.8^\circ$ per 4-h session.

The current study is not without its limitations. The prospective pilot cohort was small ($n = 127$), the clinical phase was not randomized, and the available evidence presently supports feasibility for mild-to-moderate deformity rather than definitive treatment of severe hallux valgus ($>40^\circ$). Furthermore, long-term hardware fatigue, sensor drift, maintenance burden, and production cost were not sufficiently characterized. The architecture is robust to transient connectivity loss since control is local, but dedicated network-impairment stress testing and multicentre validation are still warranted. Implementation Considerations: The most realistic near-term scenario for device implementation in terms of the potential impact on the patient is as a complement to conservative management pathways — monitored home use pre-surgical referral, symptom-guided use post-surgery, or clinician-supervised follow-up in conjunction with physiotherapy and footwear modification.

5. CONCLUSIONS

To our knowledge, the proposed wearable orthotic is the first to demonstrate that real-time patient-specific hallux valgus correction is achievable with local control updates every 100 ms, an end-to-end latency of 32.4 ± 4.7 ms, and stable execution over 219 million control steps. In the independent pilot group (127 patients), the device achieved a mean angular correction of 8.4° within 4 h, 72% short-term retention at 30 min, 41% pain reduction, greater comfort than passive bracing, and a force-tracking error of 0.237 N. These findings support the potential of adaptive conservative management for mild-to-moderate hallux valgus. Routine clinical implementation will require longer-term multicentre trials, dedicated durability and network-robustness testing, health-economic evaluation, and integration into clinical workflows for preoperative, postoperative, and non-surgical care.

REFERENCES

- [1] Cai, Y., Song, Y., He, M., et al. (2023). Global prevalence and incidence of hallux valgus: A systematic review and meta-analysis. *Journal of Foot and Ankle Research*, 16(11): 63 <https://doi.org/10.1186/s13047-023-00661-9>
- [2] Menz, H.B., Marshall, M., Thomas, M.J., Rathod-Mistry, T., Peat, G.M., Roddy, E. (2023). Incidence and progression of hallux valgus: A prospective cohort study. *Arthritis Care Res (Hoboken)*, 75(1): 166-173. <https://doi.org/10.1002/acr.24754>
- [3] Ettinger, S., Spindler, F.T., Marschall, U., Polzer, H., Stukenborg-Colsman, C., Baumbach, S.F. (2025). Hallux valgus: Prevalence and treatment options. *Deutsches Ärzteblatt International*, 122(11): 308-314. <https://doi.org/10.3238/arztebl.m2025.0068>
- [4] Roddy, E., Zhang, W., Doherty, M. (2007). Validation of a self-report instrument for assessment of hallux valgus. *Osteoarthritis Cartilage*, 15(9): 1008-1012. <https://doi.org/10.1016/j.joca.2007.02.016>
- [5] Haghi, M., Thurow, K., Stoll, R. (2017). Wearable devices in medical internet of things: Scientific research and commercially available devices. *Healthcare Informatics Research*, 23(1): 4-15. <https://doi.org/10.4258/hir.2017.23.1.4>
- [6] Quang, H.D., (2023). IoT in wearable devices: For patients and beyond. *FPT Software*. <https://fptsoftware.com/resource-center/blogs/iot-in-wearable-devices-for-patients-and-so-beyond>.
- [7] Zhang, X., Rong, X., Luo, H. (2024). Optimizing lower limb rehabilitation: The intersection of machine learning and rehabilitative robotics. *Frontiers in Rehabilitation Sciences*, 5: 1246773. <https://doi.org/10.3389/fresc.2024.1246773>
- [8] Misir, A., Yuce, A. (2025). AI in orthopedic research: A comprehensive review. *Journal of Orthopaedic Research*, 43(8): 1508-1527. <https://doi.org/10.1002/jor.26109>
- [9] Santilli, V., Mangone, M., Diko, A., et al. (2023). The use of machine learning for inferencing the effectiveness of a rehabilitation program for orthopedic and neurological patients. *International Journal of Environmental Research and Public Health*, 20(8): 5575. <https://doi.org/10.3390/ijerph20085575>
- [10] Swarnakar, R., Yadav, S.L. (2023). Artificial intelligence and machine learning in motor recovery: A rehabilitation medicine perspective. *World Journal of Clinical Cases*, 11(29): 7258-7260. <https://doi.org/10.12998/wjcc.v11.i29.7258>
- [11] Zhang, G.H., Chen, Y.D., Zhou, W.X., Chen, C.Y., Liu, Y. (2023). Bioenergy-based closed-loop medical systems for the integration of treatment, monitoring, and feedback. *Small Science*, 3(10): 2300043. <https://doi.org/10.1002/ssmc.202300043>
- [12] Paton, J.F.R., Žera, T., Vadigepalli, R., Herring, N., Pattersin, D.J. (2025). Multimodal, device-based therapeutic targeting of the cardiovascular autonomic nervous system. *Nature Reviews Cardiology*, 23: 255-278. <https://doi.org/10.1038/s41569-025-01212-4>
- [13] U.S. Food and Drug Administration. (2023). Technical considerations for medical devices with physiologic closed-loop control technology: Guidance for industry and FDA staff. Silver Spring, MD: FDA.
- [14] Shefa, F.R., Sifat, F.H., Uddin, J., Ahmad, Z., Kim, J.M., Kibria, M.G. (2024). Deep learning and IoT-based ankle-foot orthosis for enhanced gait optimization. *Healthcare*, 12(22): 2273. <https://doi.org/10.3390/healthcare12222273>
- [15] Santos, V.M., Gomes, B.B., Neto, M.A., Amaro, A.M. (2024). A systematic review of insole sensor technology: Recent studies and future directions. *Applied Sciences*, 14(14): 6085. <https://doi.org/10.3390/app14146085>
- [16] Resch, S., Kousha, A., Carroll, A., et al. (2025). Smart device development for gait monitoring: Multimodal feedback in an interactive foot orthosis, walking aid, and mobile application. *Technologies*, 13(12): 588. <https://doi.org/10.3390/technologies13120588>
- [17] Molle, R., Tamantini, C., Lauretti, C., Romano, E.M., Zollo, L. (2025). An online reinforcement learning method to improve control adaptability in robot-aided rehabilitation. *Engineering Applications of Artificial Intelligence*, 161: 112248. <https://doi.org/10.1016/j.engappai.2025.112248>
- [18] Pareek, S., Nisar, H.J., Kesavadas, T. (2024). AR3n: A reinforcement learning-based assist-as-needed controller for robotic rehabilitation. *IEEE Robotics & Automation Magazine*, 31(3): 74-82. <https://doi.org/10.1109/MRA.2023.3282434>
- [19] Pelosi, A.D., Roth, N., Yehoshua, T., et al. (2024). Personalized rehabilitation approach for reaching movement using reinforcement learning. *Scientific Reports*, 14(1): 17675. <https://doi.org/10.1038/s41598-024-64514-6>
- [20] Yang, R., Zheng, J., Song, R. (2022). Continuous mode adaptation for cable-driven rehabilitation robot using reinforcement learning. *Frontiers in Neurorobotics*, 16: 1068706. <https://doi.org/10.3389/fnbot.2022.1068706>
- [21] Kenry, Yeo, J., Lim, C. (2016). Emerging flexible and wearable physical sensing platforms for healthcare and biomedical applications. *Microsystems & Nanoengineering*, 2: 16043. <https://doi.org/10.1038/micronano.2016.43>
- [22] Vaghasiya, J.V., Mayorga-Martinez, C.C., Pumera, M. (2023). Wearable sensors for telehealth based on emerging materials and nanoarchitectonics. *npj Flexible Electronics*, 7: 26. <https://doi.org/10.1038/s41528-023-00261-4>
- [23] Deng, Z., Guo, L., Chen, X., Wu, W. (2023). Smart wearable systems for health monitoring. *Sensors*, 23(5): 2479. <https://doi.org/10.3390/s23052479>
- [24] Zheng, Z., Zhu, R., Peng, I., Xu, Z., Jiang, Y. (2024). Wearable and implantable biosensors: Mechanisms and applications in closed-loop therapeutic systems. *Journal of Materials Chemistry B*, 12: 8577-8604. <https://doi.org/10.1039/d4tb00782d>
- [25] LeCun, Y., Bengio, Y., Hinton, G. (2015). Deep learning. *Nature*, 521: 436-444. <https://doi.org/10.1038/nature14539>
- [26] Steinhubl, S.R., Muse, E.D., Topol, E.J. (2013). Can mobile health technologies transform health care? *JAMA*, 310(22): 2395-2396. <https://doi.org/10.1001/jama.2013.281078>
- [27] Banks, A., Briggs, E., Borgendale, K., Gupta, R. (2019). MQTT Version 5.0. OASIS Standard.
- [28] Amazon Web Services. (2024). AWS IoT Core Developer Guide. Amazon Web Services.
- [29] TensorFlow Lite Team. (2024). TensorFlow Lite: Machine learning for mobile and IoT devices.

- <https://www.tensorflow.org/lite>.
- [30] Goodfellow, I., Bengio, Y., Courville, A. (2016). Deep Learning. MIT Press.
- [31] Kingma, D.P., Ba, J. (2014). Adam: A method for stochastic optimization. ArXiv preprint ArXiv:1412.6980. <https://doi.org/10.48550/arXiv.1412.6980>
- [32] Bottou, L., Curtis, F.E., Nocedal, J. (2018). Optimization methods for large-scale machine learning. SIAM Review, 60(2): 223-311. <https://doi.org/10.1137/16M1080173>
- [33] Watkins, C.J.C.H., Dayan, P. (1992). Q-learning. Machine Learning, 8: 279-292. <https://doi.org/10.1007/BF00992698>
- [34] Geyer, H., Seyfarth, A., Blickhan, R. (2006). Compliant leg behaviour explains basic dynamics of walking and running. Proceedings of the Royal Society B: Biological Sciences, 273(1603): 2861-2867. <https://doi.org/10.1098/rspb.2006.3637>
- [35] Fung, Y.C. (1993). Biomechanics: Mechanical Properties of Living Tissues (2nd ed.). Springer-Verlag, New York. <https://doi.org/10.1007/978-1-4757-2257-4>
- [36] Kwan, M.Y., Yick, K.L., Yip, J., Tse, C.Y. (2021). Hallux valgus orthosis characteristics and effectiveness: A systematic review with meta-analysis. BMJ Open, 11(8): e047273. <https://doi.org/10.1136/bmjopen-2020-047273>
- [37] Patel, S., Park, H., Bonato, P., Chan, L., Rodgers, M. (2012). A review of wearable sensors and systems with application in rehabilitation. Journal of NeuroEngineering and Rehabilitation, 9: 21. <https://doi.org/10.1186/1743-0003-9-21>