

STEP-RELATIVE DEADLINE GATING FOR WIRELESS SMART-INSOLE BIOFEEDBACK

Muhammad Suleman

Alexandra Engineering University, Al Attarin, Alexandria Governorate, Egypt

Received: 17 February 2024 Revised: 29 May 2024 Accepted: 10 June 2024 Published: 20 June 2024

Correspondence: msulemanstu@gmail.com

ABSTRACT

A wireless gait-feedback command may remain technically valid after the opportunity to prepare its intended step has passed. Receiver-side deadline gating can limit such delivery, but its cost depends on cadence prediction, clock alignment and the availability of recent contacts. To quantify deadline compliance and lost cue opportunities using independently calibrated comparators, and to determine which timing assumptions most strongly affect performance. An open-loop experiment generated four bilateral contact patterns and three wireless-delay conditions. Calibration used 40 episodes per pattern; held-out evaluation used 120 episodes per pattern and link, each with 500 contacts. Nine policies were evaluated at a 150-ms preparation reserve. Comparator constants were selected on calibration data to maximize on-time coverage subject to at most 1% deadline violations among emitted cues. Paired, stratified episode bootstrap intervals quantified Monte Carlo uncertainty. Clock drift, correlated transport degradation, actuator jitter, command reordering, alternating step asymmetry and missed acquisition contacts were tested separately and jointly. The primary evaluation comprised 207,204 cue-request/link instances. Under congestion, unrestricted delivery provided 69.76% on-time coverage with 23.99% late emitted cues. A 12-interval tenth-quantile gate with a 20-ms margin provided 61.36% coverage with 0.65% late cues. Independently calibrated rolling and fixed gates provided 63.69% and 52.75% coverage, respectively, at similar late fractions of 0.93% and 0.96%. Clock offset and drift raised the reference rolling gate's late fraction to 3.28%; 5% missing acquisition contacts reduced its coverage to 31.75%. Step-relative gating offers a measurable timing advantage over a single calibrated nominal period within the evaluated models. Its benefit is conditional on clock calibration and a sufficiently complete contact history. The results establish computational timing behaviour; human preparation time, cue interpretation and clinical effectiveness remain unmeasured.

Keywords: smart insole; gait-event timing; vibrotactile feedback; deadline admission; cadence prediction; clock drift; Monte Carlo simulation.

INTRODUCTION

Pressure and inertial sensors provide contact timing and other gait measurements during walking and running (Muro-de-la-Herran et al., 2014; Camomilla et al., 2018; Kobsar et al., 2020; Mason et al., 2023). When these measurements trigger a vibration, the command must reach the actuator before the intended movement deadline. A correctly detected contact can therefore produce a late cue if processing, radio transmission or actuator startup consumes the available interval.

The movement variable and the sensory instruction determine how feedback is used during motor practice (Sigrist et al., 2013; Giggins et al., 2013). He and colleagues tested pressure-based auditory gait retraining with a wireless shoe insole during laboratory and unsupervised walking (He et al., 2019). Sheerin and colleagues evaluated haptic retraining linked to tibial acceleration in runners, including retention and transfer assessments (Sheerin et al., 2020). These applications require the feedback to refer to the movement the user is being asked to change.

Zhang and colleagues used closed-loop vibrotactile control with predicted responses to anticipatory and delayed cues during overground walking (Zhang et al., 2020). Wu and colleagues evaluated individualized auditory cue selection using a

learned cue-response relation in a Parkinson's disease feasibility study (Wu et al., 2023). Cue selection determines which instruction to send. Receiver admission determines whether a transmitted instruction still has enough time to precede its target contact.

A review of consumer gait-feedback devices identified incomplete descriptions of device performance (Mate et al., 2023). Bluetooth configuration affects transmission delay and energy use, and wireless actuator control must account for messages arriving out of order (Gomez et al., 2012; Park et al., 2019). This study evaluates a receiver rule that compares expected vibration onset with a next-contact prediction and a 150-ms preparation reserve. It counts both timely cues per request and late cues per emission, calibrates five prediction families independently, and measures the consequences of clock error, missing contacts and actuator occupancy.

CONTACT TIMING AND RECEIVER ADMISSION

BILATERAL CONTACTS AND CUE MEANING

Let t_i denote the true initial-contact time of contact i in an alternating left-right sequence. The intended target is the next *opposite-foot* contact, t_{i+1} , and $T_i = t_{i+1} - t_i$. A same-foot stride comprises two successive intervals in this definition. The generated intervals describe cadence; walking velocity cannot be inferred without step-length information.

The timing model places a contact-history controller between bilateral insole acquisition and a wireless vibrotactile receiver. The controller combines alternating left- and right-foot timestamps on a common local clock; clock calibration converts those timestamps to receiver time. Combining contact records contributes to the processing delay. A single insole detecting only same-foot contacts would instead provide stride intervals and would require a different target definition. The acquisition unit, contact-history controller and timestamp conversion are modeled in software; no physical assembly was tested.

Each request refers to preparation of the next opposite-foot step. Requests occur with probability .30 independently of the contact intervals and the admission rule. The computation evaluates whether vibration onset leaves the specified preparation reserve; it does not assign a plantar-pressure target, classify gait quality or model vibration perception. Feedback intended to change loading during the current stance phase would require a deadline tied to that phase.

RECEIPT, PREDICTED ONSET AND THE PREPARATION RESERVE

The received command carries an event identifier, foot identity, timestamp $\hat{t}_i$ and the predictor computed at contact i . Let D_i represent detection, consolidation and software delay, L_i wireless delay, and V_i actual actuator startup delay. Receipt and actual onset are

$$r_i = t_i + D_i + L_i, \quad a_i = r_i + V_i. \quad (1)$$

At receipt, the controller can compute *expected* onset $\bar{a}_i = r_i + \bar{V}$ and estimated age $g_i = \bar{a}_i - \hat{t}_i$. The primary model uses $\bar{V} = V_i = 18$ ms; jitter tests separate the predicted and actual values. A positive preparation reserve H defines

$$\begin{cases} S_i = t_{i+1} - a_i, \\ \text{on time if } S_i \geq H, \\ \text{late if } S_i < H, \\ \text{stale if } S_i \leq 0. \end{cases} \quad (2)$$

The main reserve is $H = 150$ ms; separate calculations use 100 and 200 ms. Lateness means that actuator onset leaves less than this reserve before the next contact. These values are timing allowances and have not been validated as human response thresholds. Figure 1 shows the target contact and its deadline.

CONSERVATIVE INTERVAL PREDICTORS

Observed intervals are $\tilde{T}_j = \hat{t}_{j+1} - \hat{t}_j$. At contact i , only intervals with $j < i$ are available. With a 12-interval history, the reference rolling predictor is

$$B_i^{(q)} = \max\{0.220, Q_{.10}(\tilde{T}_{i-12:i-1}) - G\}, G = 0.020 \text{ s}. \quad (3)$$

$Q_{.10}$ uses linear interpolation. The predictor uses a 20-ms subtraction and a 220-ms floor. The floor is 50 ms below the minimum generated true interval of 270 ms and limits the effect of unusually short observed intervals. History length,




FIGURE 1. TIMING RELATIVE TO A TRUE CONTACT AT 0 MS AND AN INTENDED OPPOSITE-FOOT CONTACT AT 520 MS. A 150-MS RESERVE PLACES THE DEADLINE AT 370 MS. THE LOWER MARKER REPRESENTS A POTENTIAL ACTUATOR ONSET IF THAT COMMAND WERE EMITTED; THE GATE WOULD WITHHOLD IT. NEITHER TRACE IS A MEASURED MOVEMENT OR DEVICE RECORD.

quantile and subtraction are varied in the parameter analysis. A pace-aware alternative uses medians m_4 of the latest four intervals and m_8 of the preceding eight:

$$\begin{cases} d_i = \min\{0.100, \max(0, m_8 - m_4)\}, \\ B_i^{(p)} = \max\{0.220, \min[Q_{.10}, m_4 - d_i] - G\}. \end{cases} \quad (4)$$

An available request is admitted when

$$g_i + H \leq B_i. \quad (5)$$

These quantities are conservative predictors, not distribution-free lower confidence bounds. An unexpected short interval can violate either prediction. Neither rule accesses t_{i+1} when deciding admission.

A CONDITIONAL TIMING GUARANTEE

A clock-error allowance and an actuator-startup allowance enter the admission inequality as follows. Write $e_i = \hat{t}_i - t_i$ and $j_i = V_i - \bar{V}$.

Proposition 1. *If $T_i \geq B_i$, $e_i \leq E$ and $j_i \leq J$, then admission under*

$$g_i + H + E + J \leq B_i \quad (6)$$

implies $S_i \geq H$.

Proof. By definition, $a_i - t_i = g_i + e_i + j_i$. Equation 6 therefore gives $a_i - t_i \leq B_i - H \leq T_i - H$, and hence $t_{i+1} - a_i \geq H$. $\square$

The proposition identifies the required premises; it does not establish them for the empirical predictors. Gaussian timestamp error is unbounded, and the next interval is uncertain. The primary experiment evaluates Equation 5 without additional E or J allowances. Engineering margins cannot replace measurement of clock error or behavioural preparation time.

SIMULATION METHODS

INDEPENDENT CALIBRATION AND EVALUATION

Calibration comprised 40 episodes for each of four contact patterns under the congested link, giving 160 episodes. Independent evaluation comprised 120 episodes for each pattern under each of three links. Each episode contained 500 contacts. Contacts 0–19 accumulated history, and contact 499 had no subsequent contact for deadline classification, leaving 479 eligible positions. The episode counts were fixed before execution. Episode-level uncertainty estimates follow simulation-study guidance (Morris et al., 2019).

Gait histories, requests, processing delays and nominal timestamp perturbations were paired across links. Delay and loss streams were link specific. Calibration and evaluation used disjoint random-stream families. Consequently, the primary experiment includes 480 distinct held-out gait histories exposed to three links, yielding 1,440 gait-link episodes; the repeated link exposures are not independent participant observations.

For each eligible contact, a Bernoulli draw with probability .30 generated an externally specified request. This isolates timing decisions from the ability to detect an undesirable gait feature. Feedback did not change subsequent contacts. All policies within an episode received the same requests, delays, losses and timestamp perturbations.

CONTACT AND COMMUNICATION PROCESSES

Table 1 specifies the four generated cadence patterns. An autoregressive perturbation $u_i = \rho u_{i-1} + \varepsilon_i$ used $u_0 \sim N(0, \sigma^2)$ and $\varepsilon_i \sim N[0, \sigma^2(1 - \rho^2)]$. All intervals were clipped to $[.270, .800]$ s. Values were selected to expose steady timing, gradual acceleration, state transitions and irregular shortening. They do not estimate distributions in a sampled human population.

TABLE 1. SYNTHETIC CONTACT-INTERVAL GENERATORS. UNITS ARE MILLISECONDS EXCEPT FOR PROBABILITIES AND AUTOREGRESSIVE COEFFICIENTS.

Pattern	Deterministic or switching component	Perturbation
Steady	Constant 580-ms interval.	$\sigma = 14; \rho = .55$.
Accelerating	Linear change from 640 to 350 ms over 500 interval positions.	$\sigma = 18; \rho = .55$.
Switching	Markov switching between 635 and 395 ms; .038 probability of changing state per contact.	$\sigma = 16; \rho = .55$.
Irregular	465-ms interval; independent .05 probability of an additional uniformly distributed 65–120-ms shortening.	$\sigma = 54; \rho = .40$.

Wireless delay was lognormal with median M and log-scale standard deviation s . An independent tail event added a uniform increment; independent packet loss removed a request before admission (Table 2). Processing delay was $D_i = 12 + U(0, 18)$ ms, and nominal timestamp error was independent $N(0, 8^2)$ ms. The model represents the timing consequences of heterogeneous communication; the values are not claimed as measured Bluetooth performance (Gomez et al., 2012).

TABLE 2. WIRELESS-DELAY AND PACKET-LOSS PARAMETERS. THE TAIL INCREMENT IS ADDED WITH PROBABILITY p_b ; p_ℓ IS PACKET-LOSS PROBABILITY.

Link	M (ms)	s	p_b (%)	p_ℓ (%)	Tail increment (ms)
Low jitter	40	.40	1.0	1.0	$U(100, 180)$
Mixed jitter	95	.55	8.0	3.5	$U(150, 300)$
Congested	145	.55	16.0	8.0	$U(200, 440)$

COMPARATORS AND CALIBRATION CRITERION

Unrestricted delivery admitted every available request. A nominal fixed comparator used $B_i = 450$ ms. Additional fixed, previous-interval and exponentially weighted moving-average (EWMA) comparators allowed a fairer assessment of adaptive prediction. The previous-interval predictor used $B_i = \max(.220, \tilde{T}_{i-1} - G)$. EWMA used $\mu_1 = \tilde{T}_0$ and

$$\begin{cases} \mu_i = .35\tilde{T}_{i-1} + .65\mu_{i-1}, \\ B_i^{(e)} = \max(.220, \mu_i - G). \end{cases} \quad (7)$$

The coefficient .35 was a specified comparator constant, not optimized on evaluation data.

Calibration selected a fixed period from 270–650 ms in 5-ms increments, or a subtracted margin from 0–150 ms in 5-ms increments for each adaptive family. At $H = 150$ ms, the criterion was maximum pooled on-time coverage among settings with at most 1% late emitted cues on calibration episodes. The criterion is a model-specific engineering budget, not a clinical safety limit. The reference rolling and pace-aware policies retained their 20-ms margin independently of this selection. Table 3 gives the five selected settings; none was retuned on held-out or stress-test outcomes.

TABLE 3. CALIBRATION OUTCOMES FROM 22,843 GENERATED REQUESTS UNDER THE CONGESTED LINK. COVERAGE AND LATE FRACTIONS ARE PERCENTAGES.

Policy	Selected quantity	Value (ms)	Coverage	Late / emitted
Fixed calibrated	Period	380	53.50	0.86
Previous interval	Subtracted margin	40	64.21	0.91
EWMA	Subtracted margin	50	63.67	0.93
Rolling calibrated	Subtracted margin	5	64.52	0.94
Pace calibrated	Subtracted margin	0	63.64	0.95

OUTCOMES AND MONTE CARLO INTERVALS

Let N_R be generated requests, N_I emitted cues, N_O emitted on-time cues, N_L emitted late cues and N_S emitted stale cues. The estimands were

$$C = \frac{N_O}{N_R}, \quad A = \frac{N_I}{N_R}, \quad \ell = \frac{N_L}{N_I}, \quad s_c = \frac{N_S}{N_I}. \quad (8)$$

C is on-time coverage and A is the emitted fraction. Packet loss, failed history requirements and suppression reduce coverage. Late and stale proportions are conditional on emission and are undefined if no cue is emitted. The rate N_L/N_R was also computed to expose denominator changes when comparing policies. Counts satisfy $N_I = N_O + N_L$ and $N_S \leq N_L$.

Rates were pooled over complete episodes, preserving the specified composition of four patterns. For each condition, 2,000 bootstrap replicates resampled the 120 episodes with replacement within each pattern. The same resampled episode indices were applied to every policy, giving paired contrasts. Percentile limits at 2.5 and 97.5% quantify Monte Carlo uncertainty under the stated generators. They are conditional on the selected calibration dataset and do not include uncertainty from repeatedly recalibrating the policies. They do not describe a patient population. No biological hypothesis tests were performed.

STRESS TESTS AND PARAMETER SENSITIVITY

Eight profiles used the held-out congested-link histories. The reference profile retained the primary assumptions. Each remaining profile changed the following components:

1. **Burst link:** a stationary two-state chain entered degradation with probability .04 per contact and exited with probability .22. The degraded-state probability was .154. Degraded contacts received the 200–440-ms tail increment and .35 packet-loss probability; normal-state loss was .03. This changed both serial dependence and the association between delay and loss.
2. **Clock drift:** timestamp error included a positive 30-ms offset plus $150 \times 10^{-6} t_i$ seconds of drift, in addition to the nominal 8-ms random error. The receiver retained its nominal calibration.
3. **Actuator jitter:** actual startup delay was clipped $N(18, 6^2)$ ms on $[4, 40]$ ms, while admission continued to use an expected 18 ms.
4. **Ordering and busy actuator:** .05 probability of an additional $U(350, 800)$ -ms outage increment; received commands with an identifier older than the highest received identifier were rejected. A receiver discarded arriving commands while an earlier admitted cue was starting or within its 80-ms vibration duration. No command was queued.
5. **Asymmetric steps:** alternating +50 and –50 ms changes were added before clipping the true intervals. This was a timing perturbation, not a validated model of a clinical gait disorder.
6. **Missing contacts:** each acquisition contact record was independently unavailable to the contact-history controller with probability .05. Acquisition sequence identifiers flagged omitted records. A missing current record prevented command availability, and an interval spanning an omitted record was invalid. Rolling and pace-aware policies withheld requests until a complete contiguous history was available. Previous-interval prediction required one valid adjacent interval. EWMA reset after an invalid interval and restarted from the next valid one. An undetected physical contact without an identifier flag was outside this test.
7. **Combined:** all six changes above were applied together.

Externally generated requests at contacts with missing records remained in N_R , representing unavailable cue opportunities. Loss on the acquisition path was therefore distinguished from loss on the subsequent command link. The primary three-link experiment did not include actuator occupancy or command-reordering filters; both were active only in the ordering/occupancy and combined tests.

Preparation reserves of 100, 150 and 200 ms were evaluated with identical random streams and unchanged calibrated constants. Rolling-predictor sensitivity used histories of 6, 12 and 24 intervals, quantiles .05, .10 and .20, and margins of 0, 20 and 40 ms, producing 27 combinations. For a 24-interval history, otherwise eligible contacts 20–23 could not yet be admitted; this startup cost remained in coverage.

IMPLEMENTATION AND VERIFICATION

The master seed was 20240117. NumPy SeedSequence received the tuple (master seed, stage, pattern, episode, stream, link). Stage 0 denoted calibration and stage 1 held-out evaluation; link was zero for channels shared across links. Separate stream identifiers generated contacts, timestamp noise, processing delay, requests, transport delay, loss, tail events, tail

sizes, missing contacts, degraded-link states, extra outages and actuator jitter. All channels and their ordering are specified in the complete program embedded in the LaTeX source.

Execution used Python 3.12.14, NumPy 2.3.5, pandas 2.2.3 and Matplotlib 3.10.8. Verification checked count identities, predictor invariance when future observations were changed, and the algebraic guarantee using bounded synthetic errors satisfying its premises. These checks address implementation consistency; they are not physical validation.

RESULTS

PRIMARY TIMING OUTCOMES

The held-out histories generated 69,068 requests per link, or 207,204 request/link instances across the three exposures. Each instance was evaluated by all nine policies. Table 4 shows coverage and conditional late fractions. Low-jitter unrestricted delivery already supplied 98.96% on-time coverage with 0.08% late emitted cues. Reference rolling gating supplied 98.74% and 0.01%, respectively. Under mixed jitter, unrestricted delivery supplied 89.57% coverage with 7.19% late cues, whereas the reference rolling rule supplied 84.41% coverage with 0.33% late cues.

TABLE 4. HELD-OUT TIMING OUTCOMES AT $H = 150$ MS. EACH CELL IS ON-TIME COVERAGE / LATE AMONG EMITTED CUES (%). EACH LINK HAS 69,068 REQUESTS; REFERENCE PREDICTOR MARGINS ARE 20 MS.

Policy	Low jitter	Mixed jitter	Congested
Send all	98.96 / 0.08	89.57 / 7.19	69.76 / 23.99
Fixed 450	98.96 / 0.07	85.29 / 1.94	63.10 / 4.55
Fixed calibrated	98.60 / 0.02	79.31 / 0.38	52.75 / 0.96
Previous interval	98.74 / 0.01	85.47 / 0.50	63.31 / 1.05
EWMA	98.80 / 0.01	85.21 / 0.45	62.69 / 0.97
Rolling Q10	98.74 / 0.01	84.41 / 0.33	61.36 / 0.65
Pace aware	98.48 / 0.01	82.69 / 0.28	59.51 / 0.59
Rolling calibrated	98.82 / 0.02	85.98 / 0.47	63.69 / 0.93
Pace calibrated	98.74 / 0.02	85.23 / 0.45	62.84 / 0.91

Under congestion, unrestricted delivery emitted 63,394 cues: 48,183 were on time, 15,211 were late and 6,467 were stale. The reference rolling rule emitted 42,657 cues, including 42,380 on-time, 277 late and 11 stale cues. Its coverage was 61.36 [60.97, 61.74]%, and its conditional late fraction was 0.65 [0.57, 0.73]%

Relative to unrestricted delivery, reference rolling gating reduced the conditional late fraction by 23.35 percentage points (95% Monte Carlo interval 23.00–23.69), while reducing on-time coverage by 8.40 points (8.20–8.61). It suppressed 14,934 late cues and 5,803 cues that would have met the deadline. The latter constituted 12.04% of the unrestricted policy's timely cues. Expressed against all requests, late delivery fell from 22.02% to 0.40%; thus the improvement was not produced solely by changing the conditional denominator.

The reference pace-aware rule provided 59.51% coverage with 0.59% late cues. Compared with reference rolling gating, it forfeited a further 1.85 coverage points (1.76–1.96) for a 0.064-point reduction in conditional lateness (0.038–0.094). The maximum half-width of the primary 95% intervals for coverage or lateness across all policies and links was 0.394 percentage points.

CALIBRATION AND COMPARATOR PERFORMANCE

The nominal 450-ms fixed rule produced 63.10% coverage under congestion, but 4.55% of emitted cues were late. Calibration reduced the selected fixed period to 380 ms. Its held-out coverage was 52.75% with 0.96% late cues. Rolling calibration selected a 5-ms margin and provided 63.69% coverage with 0.93% late cues. The rolling-minus-fixed coverage difference was 10.94 points (10.65–11.23), while the late-fraction difference was –0.028 points (–0.131 to 0.079). The calibrated rolling rule therefore retained 7,556 more on-time cues than the calibrated fixed rule, while its late fraction differed by less than 0.03 percentage points.

The calibrated previous-interval policy provided 63.31% coverage with 1.05% late cues. Its held-out point estimate exceeded the calibration criterion, illustrating that a training budget is not a deployment guarantee. Calibrated EWMA provided 62.69% coverage with 0.97% late cues; calibrated pace-aware gating provided 62.84% with 0.91%. The three calibrated adaptive rules had coverage within 0.63 percentage points and late fractions between 0.91% and 1.05%. Figure 2 displays held-out parameter sweeps within the congested condition; the calibration-selected points were fixed independently of these curves.




FIGURE 2. HELD-OUT COVERAGE-LATENCY CURVES FOR CANDIDATE FIXED PERIODS AND ADAPTIVE MARGINS UNDER CONGESTION. CIRCLES IDENTIFY INDEPENDENTLY CALIBRATED SETTINGS. THE DASHED LINE IS THE 1% CALIBRATION CRITERION, NOT A SAFETY THRESHOLD. THE VERTICAL DISPLAY IS RESTRICTED TO 0–6% TO DISTINGUISH THE CANDIDATE GATES; UNRESTRICTED DELIVERY HAS 23.99% LATENCY AND APPEARS IN TABLE 4. FULL PARAMETER SWEEPS ARE GENERATED BY THE EMBEDDED PROGRAM.

CLOCK ALIGNMENT AND HISTORY AVAILABILITY

Figure 3 and Table 5 summarize the eight profiles. Adding a 30-ms positive clock offset and 150-ppm drift left true unrestricted timing unchanged. It increased reference rolling coverage to 67.20%, but increased its late fraction to 3.28%. Apparent age was underestimated, admitting commands the correctly calibrated gate would have withheld. EWMA's late fraction similarly rose to 3.09%.

Missing acquisition records prevented admission: the requirement for a complete 12-interval history reduced reference rolling coverage from 61.36% to 31.75%. Its late fraction remained 0.60%, conditional on the smaller emitted set. EWMA, which restarted after a single restored adjacent interval, retained 56.48% coverage with 0.96% latency. A low conditional late fraction therefore concealed substantial loss of cue availability when the rolling history was repeatedly invalidated.

Alternating step asymmetry reduced rolling coverage to 54.08%, with 1.01% late cues. EWMA retained 61.27% coverage but had 3.21% latency. The joint profile yielded 32.14% coverage and 3.70% latency for rolling gating, compared with 59.14% and 6.84% for EWMA. These results show different sensitivity to complete-history requirements and interval smoothing; neither policy met the calibration budget throughout the stress tests.

TABLE 5. HELD-OUT STRESS OUTCOMES UNDER THE CONGESTED BASE LINK, WITH $H = 150$ MS AND UNCHANGED POLICY CONSTANTS. COVERAGE IS RELATIVE TO ALL REQUESTS; LATE FRACTIONS ARE RELATIVE TO EMITTED CUES.

Profile	Rolling Q10, 20-ms margin		Calibrated EWMA	
	Coverage (%)	Late (%)	Coverage (%)	Late (%)
Reference	61.36	0.65	62.69	0.97
Burst link	64.96	0.59	66.38	0.92
Clock drift	67.20	3.28	68.30	3.09
Actuator jitter	61.35	0.66	62.68	0.98
Ordering + busy	58.29	0.64	59.56	0.96
Asymmetric steps	54.08	1.01	61.27	3.21
Missing contacts	31.75	0.60	56.48	0.96
Combined	32.14	3.70	59.14	6.84

The burst-link profile increased rolling coverage to 64.96% with 0.59% latency. This profile coupled longer delays with greater loss probability, preferentially removing delayed requests before emission. Its difference from the reference cannot be attributed to temporal autocorrelation alone. Actuator jitter had a smaller effect: reference rolling coverage was 61.35% and latency 0.66%. Ordering rejection and actuator occupancy reduced coverage to 58.29%, with 0.64% latency; this loss followed command unavailability and occupancy suppression without changing the prediction margin.

PREPARATION RESERVE AND PREDICTOR SENSITIVITY

For $H = 100, 150$ and 200 ms, reference rolling coverage under congestion was 69.12%, 61.36% and 50.09%, respectively, with late fractions of 0.38%, 0.65% and 1.17%. Unrestricted coverage was 75.19%, 69.76% and 62.50%, with 18.08%, 23.99%




FIGURE 3. COVERAGE AND CONDITIONAL LATENESS ACROSS THE EIGHT HELD-OUT PROFILES. POINTS ARE POOLED ESTIMATES AND BARS ARE PAIRED, STRATIFIED 95% EPISODE-BOOTSTRAP MONTE CARLO INTERVALS. SOME INTERVALS ARE NARROWER THAN THE MARKERS AT THIS SCALE. THE VERTICAL DASHED LINE MARKS 1% LATENESS FOR REFERENCE. "ORDERING + BUSY" COMBINES EXTRA OUTAGE DELAY, IDENTIFIER-BASED REJECTION AND AN OCCUPIED-ACTUATOR DISCARD RULE.

and 31.91% lateness. The reserve therefore changes both the set of acceptable cues and the cost of admission; the experiment does not select a biologically preferable reserve.

Across the 27 predictor combinations, coverage ranged from 52.96% to 66.38%, and conditional lateness ranged from 0.32% to 1.70%. With 12 intervals and the tenth quantile, changing the margin from 0 to 20 to 40 ms changed coverage from 64.41% to 61.36% to 57.33%, and lateness from 1.10% to 0.65% to 0.48%. Longer histories and smaller quantiles usually reduced exposure at the cost of a more restrictive admission estimate (Figure 4).

The actual next interval fell below the reference rolling prediction in 0.46%, 2.09%, 4.53% and 9.97% of requested positions for steady, accelerating, switching and irregular patterns, respectively. The corresponding pace-aware values were 0.28%, 1.31%, 3.70% and 8.30%. These prediction breaches were more frequent than emitted late cues because some breached predictions were associated with requests that were lost or withheld, or with arrivals early enough to remain on time.




FIGURE 4. HELD-OUT ROLLING-PREDICTOR SENSITIVITY UNDER CONGESTION AT A 150-MS RESERVE. COLOUR AND MARKER DISTINGUISH HISTORY LENGTH; LINE PATTERN DISTINGUISHES THE SUBTRACTED MARGIN. EACH PLOTTED SETTING USES THE SAME GENERATED REQUESTS AND TIMING CHANNELS. THESE ARE PARAMETER DIAGNOSTICS, NOT SETTINGS SELECTED RETROSPECTIVELY FOR THE PRIMARY COMPARISON.

DISCUSSION

CUE ADMISSION AND FOREGONE INSTRUCTION

Admission tests the expected onset of each received request against a prediction of its next contact. Under congestion, the 20-ms rolling gate removed 14,934 late cues while also withholding 5,803 cues that unrestricted delivery would have emitted on time. Coverage and late fraction quantify these two consequences separately. The fixed 450-ms gate retained 63.10% coverage but produced 4.55% late emissions. Selecting 380 ms on the independent calibration episodes reduced

its late fraction to 0.96% and its coverage to 52.75%; the calibrated rolling gate retained 63.69% coverage at 0.93% lateness.

Receiver admission operates after the pressure-based, haptic or adaptive cue-generation procedures examined in gait-feedback studies (He et al., 2019; Sheerin et al., 2020; Zhang et al., 2020; Wu et al., 2023). Its purpose is to prevent an instruction from being delivered after its allotted preparation interval. The calibrated previous-interval, EWMA and pace-aware rules retained 63.31%, 62.69% and 62.84% coverage, respectively. Their similar primary timing performance makes clock tolerance and recovery after missing records relevant to selecting among the predictors.

CLOCK ERROR AND MISSING CONTACTS REQUIRE DIFFERENT RESPONSES

Clock misalignment changes estimated command age while leaving true movement and delivery time unchanged. A receiver may consequently appear to increase cue exposure even as its deadline compliance deteriorates. The conditional guarantee makes the role of bounded clock and startup uncertainty explicit, but empirical interval prediction remains an additional uncertain premise. A physical system needs calibration-health monitoring and an admissible-error policy before its nominal margin can be interpreted as reliable protection against mistiming.

A missing acquisition record removes the adjacent intervals needed for prediction. With .05 omission probability, requiring 12 contiguous intervals reduced rolling coverage to 31.75%, although only 0.60% of emitted cues were late. EWMA restarted after one valid adjacent interval and retained 56.48% coverage with 0.96% lateness. An implementation must distinguish missing acquisition records, transport loss, invalid history, deadline rejection, older identifiers and actuator occupancy. An observed interval that spans an omitted record must be excluded from single-step prediction. This calculation assumes that sequence identifiers reveal omitted records; it does not establish detection of an unflagged missed physical contact.

The joint test reduced rolling coverage to 32.14% with 3.70% lateness, while EWMA retained 59.14% coverage with 6.84% lateness. Faster history recovery therefore preserved more instruction but did not resolve clock-induced admission errors. Selecting a recovery rule requires measured contact-acquisition reliability and a specified minimum number of usable cues for the task.

KINESIOLOGICAL INTERPRETATION

The admission test requires onset at least 150 ms before the next opposite-foot contact. It does not measure perception, kinematics or learning. Reviews of rehabilitation feedback emphasize that behavioural effects depend on population, target variable and intervention characteristics (Bowman et al., 2021; Escamilla-Nunez et al., 2020). Feedback aimed at medializing plantar pressure during stance, for example, has a different target from the next-step preparation reserve in this experiment. Likewise, a runner's contact sequence or a neurologically impaired gait requires separate timing and cue-content calibration.

The numerical cost also has a training interpretation: under reference congestion, the rolling gate withheld approximately one eighth of the cues that unrestricted delivery would have supplied on time. Whether the resulting dose remains useful is unknown. Increasing radio reliability, relocating processing or locating the actuator on a shared-clock device may preserve opportunity more effectively than increasing the discard margin, but each design still needs a measured end-to-end timing budget.

CONCLUSION

Step-relative admission reduced deadline violations in a synthetic wireless smart-insole feedback system while sacrificing a measurable proportion of timely cue opportunities. Independent calibration showed that adaptive prediction retained more on-time coverage than a single fixed period at similar lateness under the evaluated congested link. Clock drift and missing acquisition contacts exposed distinct weaknesses: permissive admission through underestimated age, and loss of coverage through invalid history. A usable physical implementation therefore needs verified clock alignment, explicit history-recovery behaviour and task-specific evaluation of the preparation reserve. The computational results support these engineering requirements without establishing a clinical or sporting benefit.

DECLARATIONS

CONFLICT OF INTEREST

The author declares that there are no conflicts of interest related to this study.

DATA AVAILABILITY

The data supporting the findings of this study are available from the corresponding author upon reasonable request.

FUNDING

This research received no external funding.

REFERENCES

1. Bowman T, Gervasoni E, Arienti C, et al. Wearable devices for biofeedback rehabilitation: a systematic review and meta-analysis to design application rules and estimate the effectiveness on balance and gait outcomes in neurological diseases. *Sensors*. 2021;21(10):3444. doi:10.3390/s21103444.
2. Camomilla V, Bergamini E, Fantozzi S, Vannozzi G. Trends supporting the in-field use of wearable inertial sensors for sport performance evaluation: a systematic review. *Sensors*. 2018;18(3):873. doi:10.3390/s18030873.
3. Escamilla-Nunez R, Michelini A, Andrysek J. Biofeedback systems for gait rehabilitation of individuals with lower-limb amputation: a systematic review. *Sensors*. 2020;20(6):1628. doi:10.3390/s20061628.
4. Giggins OM, Persson UM, Caulfield B. Biofeedback in rehabilitation. *J Neuroeng Rehabil*. 2013;10:60. doi:10.1186/1743-0003-10-60.
5. Gomez C, Oller J, Paradells J. Overview and evaluation of Bluetooth Low Energy: an emerging low-power wireless technology. *Sensors*. 2012;12(9):11734–11753. doi:10.3390/s120911734.
6. He J, Lippmann K, Shakoor N, Ferrigno C, Wimmer MA. Unsupervised gait retraining using a wireless pressure-detecting shoe insole. *Gait Posture*. 2019;70:408–413. doi:10.1016/j.gaitpost.2019.03.021.
7. Hoffmann TC, Glasziou PP, Boutron I, et al. Better reporting of interventions: template for intervention description and replication (TIDieR) checklist and guide. *BMJ*. 2014;348:g1687. doi:10.1136/bmj.g1687.
8. Kobsar D, Charlton JM, Tse CTF, et al. Validity and reliability of wearable inertial sensors in healthy adult walking: a systematic review and meta-analysis. *J Neuroeng Rehabil*. 2020;17:62. doi:10.1186/s12984-020-00685-3.
9. Mason R, Pearson LT, Barry G, et al. Wearables for running gait analysis: a systematic review. *Sports Med*. 2023;53:241–268. doi:10.1007/s40279-022-01760-6.
10. Mate KKV, Abou-Sharkh A, Mansoubi M, et al. Evidence for the efficacy of commercially available wearable biofeedback gait devices: consumer-centered review. *JMIR Rehabil Assist Technol*. 2023;10:e40680. doi:10.2196/40680.
11. Morris TP, White IR, Crowther MJ. Using simulation studies to evaluate statistical methods. *Stat Med*. 2019;38(11):2074–2102. doi:10.1002/sim.8086.
12. Muro-de-la-Herran A, Garcia-Zapirain B, Mendez-Zorrilla A. Gait analysis methods: an overview of wearable and non-wearable systems, highlighting clinical applications. *Sensors*. 2014;14(2):3362–3394. doi:10.3390/s140203362.
13. Park B, Nah J, Choi JY, Yoon IJ, Park P. Robust wireless sensor and actuator networks for networked control systems. *Sensors*. 2019;19(7):1535. doi:10.3390/s19071535.
14. Sheerin KR, Reid D, Taylor D, Besier TF. The effectiveness of real-time haptic feedback gait retraining for reducing resultant tibial acceleration with runners. *Phys Ther Sport*. 2020;43:173–180. doi:10.1016/j.ptsp.2020.03.001.
15. Sigrist R, Rauter G, Riener R, Wolf P. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: a review. *Psychon Bull Rev*. 2013;20(1):21–53. doi:10.3758/s13423-012-0333-8.
16. Wu TLY, Murphy A, Chen C, Kulić D. Adaptive cueing strategy for gait modification: a case study using auditory cues. *Front Neurobot*. 2023;17:1127033. doi:10.3389/fnbot.2023.1127033.
17. Zhang H, Yin Y, Chen Z, Zhang Y, Rao AK, Guo Y, Zanotto D. Wearable biofeedback system to induce desired walking speed in overground gait training. *Sensors*. 2020;20(14):4002. doi:10.3390/s20144002.