Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Curvature-Based Reference Governor with Terminal-Braking Reserves for a Virtual Gait Platform

Jinqiu Dai

2704744953@qq.com

Tianjin University of Science and Technology

Xinting Wang

Tianjin University of Science and Technology

Zeyuan Niu

Tianjin University of Science and Technology

YuXuan Zhao

Tianjin University of Science and Technology

You Jiang

Tianjin University of Science and Technology

JingYi Zuo

Tianjin University of Science and Technology

Research Article

Keywords: Reference governor, Actuator constraints, Path parameterization, Terminal braking, Time-stamped feedback

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11724155/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A Curvature-Based Reference Governor with Terminal-Braking Reserves for a Virtual Gait Platform

Abstract

Virtual gait platforms must follow changing pose requests without exhausting the actuator stroke needed to stop. For platforms with an adaptive rotation center, this requires converting nonlinear actuator geometry into a motion limit that also accounts for response uncertainty. We propose a scalar reference governor that bounds actuator-length slopes and curvature along each requested pose path, then uses these bounds to calculate speed and acceleration envelopes. The governor admits a path fraction after reserving stroke for initial reconstruction error, response mismatch, and terminal braking. Under the declared virtual actuator model, bounded braking returns the actuators to rest, allowing successive certified segments or a stationary hold to preserve actuator-constraint admissibility. Across 540 shared paths with identical complete reserves and a 0.1-mm initial offset, Curvature increased mean certified progress by 7.69% relative to the axis-aware norm baseline AxisConservative. In twenty Unity closed-loop executions with scheduled patient-state packet delays of 0–80 ms, roll tracking root-mean-square error (RMSE) was 2.62%–3.70% lower than AxisConservative across ten configuration-matched paired comparisons. Decision-level analysis links pitch attenuation to the acceleration certificate and fixed segment timing despite unused speed and stroke margins. The resulting governor combines geometric motion limits and stopping reserves to increase allowable motion within the modeled actuator constraints.

Keywords: Reference governor; Actuator constraints; Path parameterization; Terminal braking; Time-stamped feedback

1 Introduction

Adaptive rotation-center selection changes the actuator motion required to reproduce a requested terrain tilt on a virtual gait platform. Earlier moving-platform studies established methods for changing rotation axes and characterized platform response [1, 2]. On a four-actuator platform, each pose request defines a family of actuator-length trajectories. Coordinated motion must retain stroke, speed, and acceleration margins through both the requested movement and subsequent stopping.

Actuator-aware motion-cueing methods use platform kinematics and predictive constraints to regulate admissible movement [3–5]. Reference and command governors modify requested commands to satisfy state and input constraints [6, 10]. Reachability-based time-optimal path parameterization and recent Stewart-platform trajectory scaling address path-speed allocation [7, 12], while network reference governors extend constraint management to variable delays [11]. For a gait platform with an adaptive rotation center, the governor must determine how much of a requested pose path can be executed. The available stroke must also accommodate reconstruction error, uncertain actuator response, and the subsequent stop.

A fixed-axis pose path defines four nonlinear actuator-length functions. At a segment boundary, lengths reconstructed from the platform pose can differ from the independently integrated actuator lengths. Response-time mismatch can also leave residual actuator velocity after the desired path has ended. A usable motion limit must account for both discrepancies when converting geometric path bounds into admissible actuator motion. This is a model-specific construction within the state, input, disturbance, and uncertainty framework of reference governors [6, 10].

The proposed governor uses initial directional slopes and curvature bounds to bound actuator length, speed, and acceleration along a candidate path fraction. Reconstruction error, response uncertainty, and terminal

braking determine the stroke reserve. Scalar prefix tests select a fraction that passes a common four-actuator recheck. A bounded brake returns the actuators to rest, allowing successive certified segments or a stationary hold to preserve actuator-constraint admissibility under the stated model. The contribution connects directional geometry and response-error bounds to a finite-stop supervisor; Table 1 compares its computation and guarantee with established formulations.

Gilbert and Kolmanovsky [10] govern a general nonlinear closed-loop state with bounded disturbances. An equilibrium-parameterized safety function represents admissible state deviations, and safe, strongly returnable sublevel sets justify retaining the previous reference when no advance is admitted. Online calculation selects a scalar reference increment through that function. In the present construction, four actuator lengths and velocities define the state. The measured geometry-to-length discrepancy enters an explicit response envelope, and a bounded finite brake followed by rest supplies the fallback action.

Pham and Pham's TOPP-RA [12] uses squared path speed as the discrete state and path acceleration as the control. Backward propagation computes controllable sets from a prescribed terminal-speed set; forward control selection then constructs the time parameterization. Initial speed and parameter uncertainty enter the path constraints. Here, a fixed-duration smoothstep defines the time law, and the online variable is the fraction of the requested pose path. Its derivative bounds are checked together with initial length/speed error, uncertain velocity response, and terminal-braking displacement. Proposition 1 and Corollary 1 establish actuator admissibility for this construction.

The evaluation separates the geometric contribution from its effect in the coupled platform model. Shared path and state inputs measure certified progress and decision time against the axis-aware norm baseline AxisConservative and subinterval-certified DirectSearch. Boundary responses test the error and stopping reserves. Configuration-matched Unity executions then compare tracking and motion allocation, including Curvature versus AxisConservative. The platform geometry, adaptive center-of-rotation (ACOR) controller, patient profiles, and Standard terrain are inherited from the related Chinese study [8] and specified in Appendix A.

2 Pose-path governor and actuator-state supervisor

Table 1 positions the contribution within established constraint-management mechanisms.

Table 1. Mechanism-level positioning of the model-specific construction. The established-mechanism column summarizes the mechanism families specified in its entries. Platform and ACOR provenance are given in the Introduction.

Aspect	Established mechanisms [6, 10–12]	Present construction
State model	Surveyed governors cover linear/nonlinear pre-stabilized systems; path parameterization imposes path-dependent motion constraints	Four bounded first-order velocity actuators driven by a rigid heave–pitch–roll pose prefix
Initial discrepancy	Uncertainty and disturbances enter admissibility or prediction mechanisms	Measured length discrepancy and initial speed error enter explicit scalar response envelopes
Stopping condition	Terminal, invariant-set, or backup mechanisms depend on the chosen governor	A bounded opposite-sign command reaches zero velocity; an explicit stroke reserve encloses the finite stop
Calculation	Admissible-set tests, prediction, or reachability calculations vary by method	Analytic joint-length slope/curvature envelopes and scalar bisection, followed by a four-actuator common-prefix recheck
Guarantee scope	Guarantees depend on each method's plant, uncertainty, and terminal assumptions	Recursive actuator-state admissibility under the declared response model; pose compatibility is checked separately
Increment tested	Established methods instantiate different admissibility mechanisms	The certificate and reserve construction is compared with two alternative prefix certificates under the same supervisor

2.1 Reduced platform geometry

The virtual platform is represented by the reduced pose $\mathbf{x} = (h, \phi, \psi)$, where h, ϕ, and ψ denote heave, pitch, and roll, respectively. For actuator $i \in \{1, \dots, 4\}$, let $\mathbf{b}_i$ denote its base attachment in world coordinates. Let $\mathbf{q}_i$ denote its moving-platform attachment relative to the selected rotation center. The Unity scene supplies the attachment geometry used in the model. At each resting update, heave and angular residuals are normalized by 1 m and 1 rad, respectively. The coordinate with the largest absolute normalized residual is selected; the remaining requested coordinates are deferred. These fixed scales reproduce the selection rule in canonical SI units. The selected rotation center remains fixed throughout the candidate path.

For a translation increment $\Delta\mathbf{p}$, unit rotation axis $\mathbf{a}$, signed angle θ, and path progress $s \in [0,1]$, define

$$\begin{aligned} \mathbf{p}(s) &= \mathbf{p}_0 + s\Delta\mathbf{p}, \qquad \mathbf{R}(s) = \exp(s\theta[\mathbf{a}]_\times), \\ \mathbf{r}_i(s) &= \mathbf{r}_{i,0} + s\Delta\mathbf{p} + [\mathbf{R}(s) - \mathbf{I}]\mathbf{q}_i, \\ \ell_i(s) &= \| \mathbf{r}_i(s) \|. \end{aligned} \tag{1}$$

Here $\mathbf{r}_i(s)$ is the base-to-platform attachment vector, and $\ell_i(s)$ is the corresponding virtual actuator length. The nominated rigid-platform interpolation determines all four actuator lengths, which are not planned independently.

Realized platform orientation and the consumed digital-patient center-of-mass and angular states enter the inherited ACOR controller, which updates the nominated rotation center. The terrain pose and that center define the candidate rigid path. Coordinate selection, optional information-age scaling, and the actuator supervisor determine its applied prefix. Integrated actuator states are fitted to a platform pose, which drives the digital-patient model and the next timestamped feedback updates. This information flow distinguishes the adaptive center selection from the added actuator-admissibility supervisor.

2.2 Analytic joint-length envelopes

Differentiation of (1) gives

$$\begin{aligned} \mathbf{r}_i'(s) &= \Delta\mathbf{p} + \theta(\mathbf{a} \times \mathbf{R}(s)\mathbf{q}_i), \\ \mathbf{r}_i''(s) &= \theta^2 \mathbf{a} \times (\mathbf{a} \times \mathbf{R}(s)\mathbf{q}_i). \end{aligned} \qquad (2)$$

For positive $\ell_i(s)$,

$$\begin{aligned} \ell_i'(s) &= \frac{\mathbf{r}_i^T \mathbf{r}_i'}{\ell_i}, \\ \ell_i''(s) &= \frac{(\mathbf{r}_i')^T \mathbf{r}_i' + \mathbf{r}_i^T \mathbf{r}_i'' - [\ell_i'(s)]^2}{\ell_i}. \end{aligned} \qquad (3)$$

Primes denote derivatives with respect to path progress; overdots denote derivatives with respect to time.

The path-wide envelopes are

$$B_i = \| \Delta\mathbf{p} \| + |\theta| \, \| \mathbf{q}_i \|, \qquad D_i = \theta^2 \, \| \mathbf{q}_i \|. \qquad (4)$$

They imply $|\ell_i'| \le B_i$. If $\underline{\ell}_i > 0$ is a lower bound on length over a path segment, then

$$|\ell_i''| \le C_i = \frac{B_i^2}{\underline{\ell}_i} + D_i. \qquad (5)$$

An equally inexpensive baseline removes the attachment component parallel to the unit rotation axis:

$$\begin{aligned} \mathbf{q}_{i,\perp} &= \mathbf{q}_i - \mathbf{a}(\mathbf{a}^T \mathbf{q}_i), \\ B_i^{\perp} &= \| \Delta\mathbf{p} \| + |\theta| \, \| \mathbf{a} \times \mathbf{q}_i \|, \\ D_i^{\perp} &= \theta^2 \, \| \mathbf{a} \times \mathbf{q}_i \|. \end{aligned}$$

Rotation about the fixed axis preserves the perpendicular radius. Thus the same triangle and length-derivative inequalities hold with $B_i^{\perp} \le B_i$ and $D_i^{\perp} \le D_i$. AxisConservative substitutes these bounds throughout Conservative, including its denominator, stroke excursion, and acceleration enclosure. AxisCurvature makes the identical substitution throughout Curvature. Projection is a standard geometric tightening; the comparison isolates the additional directional-slope and curvature construction after that tightening.

The direction-independent Conservative certificate bounds the slope by B_i and the stroke excursion by $B_i\alpha$ for progress α. The Curvature certificate uses the initial directional slope $k_i = |\ell_i'(0)|$. It bounds the slope by $k_i + C_i\alpha$ and the length excursion by

$$|\ell_i(s) - \ell_i(0)| \le k_i\alpha + 1/2 \, C_i\alpha^2, \qquad s \in [0, \alpha]. \qquad (6)$$

Both envelopes are evaluated for every actuator from the current feasible segment start.

For both analytic certificates, the implemented denominator bound is $\underline{\ell}_i(\alpha) = \ell_i(0) - B_i\alpha > 0$. Thus, C_i is recomputed for each tested prefix rather than treated as a fixed local curvature. Candidates with a nonpositive denominator bound are rejected before further evaluation.

2.3 Smooth path timing and candidate progress

Each segment uses a cubic smoothstep

$$\begin{aligned} s(t) &= \alpha f(t/H), \quad f(q) = 3q^2 - 2q^3, \\ q &= t/H, \qquad\quad H = 0.8 \text{ s}. \end{aligned} \tag{7}$$

The desired progress rate is zero at the beginning and end of the segment, with

$$\max|\dot{s}| = \frac{1.5\alpha}{H}, \qquad \max|\ddot{s}| = \frac{6\alpha}{H^2}. \tag{8}$$

Given path bounds $L_{1,i} \geq \sup|\ell_i'|$ and $L_{2,i} \geq \sup|\ell_i''|$, the resulting desired actuator envelopes are

$$\begin{aligned} |\dot{\ell}_i| &\leq L_{1,i}\frac{1.5\alpha}{H}, \\ |\ddot{\ell}_i| &\leq L_{2,i}\left(\frac{1.5\alpha}{H}\right)^2 + L_{1,i}\frac{6\alpha}{H^2}. \end{aligned} \tag{9}$$

Each method first checks the zero-progress candidate and tests full progress if the starting state is admissible. Bisection is then applied as needed, using 32 updates for Conservative and Curvature and 20 for DirectSearch. All actuator-specific candidates undergo the same stroke, speed, acceleration, and response-mismatch checks. The smallest candidate fraction is rechecked for all four actuators. If this common prefix fails, its fraction is repeatedly halved until a certified prefix or the resting hold is selected.

The closed-loop DirectSearch comparator uses 16 equal subintervals; the offline scan uses 1, 2, 4, 8, and 16. AxisDirectSearch applies the same search with projected-radius bounds. Each comparator divides a candidate prefix into its specified number of equal subintervals. The path Lipschitz bound B_i encloses lengths between interval nodes, while a local second-derivative envelope bounds the slope within each subinterval. Twenty bisection updates and an explicit candidate recheck return a subinterval-certified feasible fraction. The guarantee is feasibility under this procedure; global maximality is outside its search specification.

Specifically, for subinterval count N_s, let $\eta = \alpha/N_s$, $s_j = j\eta$, and $m_j = (j - 1/2)\eta$. On subinterval j, a positive midpoint-based bound $\underline{\ell}_{ij} = \ell_i(m_j) - B_i\eta/2$ gives $C_{ij} = B_i^2/\underline{\ell}_{ij} + D_i$. Let $K_{ij} = \max\{|\ell_i'(s_{j-1})|, |\ell_i'(s_j)|\}$. The search uses

$$\begin{aligned} L_{1,i} &= \max_{j=1,\ldots,N_s}\{K_{ij} + C_{ij}\eta/2\}, \\ L_{2,i} &= \max_j C_{ij}. \end{aligned}$$

Its length enclosure is $[\min_j \ell_i(s_j) - B_i\eta/2, \max_j \ell_i(s_j) + B_i\eta/2]$, with node extrema taken over $j = 0, \ldots, N_s$. These subinterval bounds certify the continuous path between nodes, rather than checking endpoints alone.

2.4 Virtual actuator response and initial-state envelope

Each actuator has a length state ℓ_i^a and velocity state v_i^a. During a commanded path segment,

$$\begin{aligned} \dot{\ell}_i^a &= v_i^a, \\ \dot{v}_i^a &= \mathrm{sat}_{[-A_{\max}, A_{\max}]}\left(\frac{u_i - v_i^a}{\tau_i}\right), \\ u_i &= \dot{\ell}_{d,i} + \tau_0 \ddot{\ell}_{d,i}. \end{aligned}$$

The declared response-time set is $\tau_i \in [\tau_{\min}, \tau_{\max}]$, with $|\tau_i - \tau_0| \le \Delta\tau$. Stroke, command, speed, acceleration, and response-time limits are specified as virtual-model assumptions. The reduced rigid-platform pose is fitted to the four integrated actuator lengths. The fitting residual is carried forward as an initial-state discrepancy for the next segment.

Let $E_{\ell,0}$ bound $|\ell_i^a(0) - \ell_{d,i}(0)|$ across the four actuators, and let $E_{v,0}$ bound $|v_i^a(0)|$. The desired actuator speed is zero at both ends of each segment. If $|\dot{\ell}_{d,i}| \le S_d$ and $|\ddot{\ell}_{d,i}| \le A_d$, define

$$E_v = E_{v,0} + \Delta\tau A_d.$$

Before saturation, $e_i = v_i^a - \dot{\ell}_{d,i}$ obeys

$$\dot{e}_i = -e_i/\tau_i + (\tau_0 - \tau_i)\ddot{\ell}_{d,i}/\tau_i.$$

The convolution solution bounds both the transient initial error and response-mismatch forcing, yielding $|e_i(t)| \le E_v$. Consequently,

$$\begin{aligned} |v_i^a| &\le S_d + E_v, \\ |u_i| &\le S_d + \tau_0 A_d, \\ |\dot{v}_i^a| &\le \left(1 + \frac{2\Delta\tau}{\tau_{\min}}\right) A_d + \frac{E_{v,0}}{\tau_{\min}}, \\ |\ell_i^a(t) - \ell_{d,i}(t)| &\le E_{\ell,0} + E_v H. \end{aligned}$$

The desired-acceleration budget is therefore

$$A_d = \frac{\max\{0, A_{\max} - E_{v,0}/\tau_{\min}\}}{1 + 2\Delta\tau/\tau_{\min}}.$$

This budget is admissible only if $E_{v,0} \le A_{\max}\tau_{\min}$. If $E_{v,0} > A_{\max}\tau_{\min}$, no nonnegative A_d satisfies the acceleration inequality in Proposition 1; clipping the numerator to zero does not certify that initial state. At equality the budget is zero and the other proposition conditions still apply. A nonzero initial velocity also requires its speed and stopping reserves to be certified; zero requested progress is a resting hold only when the actuator is already at rest. The experiments initiate segments from the resting set, with $E_{v,0} = 0$.

The governor evaluates initial errors from the integrated actuator states and reconstructed segment-start geometry. It checks both command and realized-speed envelopes against $V_{\max}$, giving a state-based speed certificate throughout the nominated motion.

2.5 Terminal braking reserve

At the path endpoint, the desired velocity is zero but response mismatch can leave $|v_i^a(H)| \le E_v$. The supervisor applies the bounded braking command

$$u_{b,i} = \begin{cases} -b\,\mathrm{sign}(v_i^a), & v_i^a \neq 0, \\ 0, & v_i^a = 0, \end{cases} \qquad 0 < b \le V_{\max},$$

through the same first-order response with acceleration saturation. While the actuator is moving, the braking deceleration has the lower bound

$$a_b = \min\{A_{\max}, b/\tau_{\max}\} > 0.$$

Speed therefore decreases monotonically, stopping occurs within E_v/a_b, and the additional displacement is bounded by $E_v^2/(2a_b)$. The complete stroke reserve for the path is

$$\delta_\ell = E_{\ell,0} + E_v H + \frac{E_v^2}{2a_b}.$$

The virtual-model parameters are $[\ell_{\min}, \ell_{\max}] = [0.380, 0.630]$ m, $H = 0.8$ s, $V_{\max} = 0.0095$ m/s, $A_{\max} = 0.050$ m/s^2, $\tau_0 = 0.060$ s, $\Delta\tau = 0.030$ s, and $[\tau_{\min}, \tau_{\max}] = [0.030, 0.090]$ s. These are declared simulation assumptions. With $b = 0.0015$ m/s, the resulting braking bound is $a_b = 0.01667$ m/s^2. For a resting start, $A_d = 0.01667$ m/s^2, and the combined response and braking reserve is 0.4075 mm. The measured initial length discrepancy is added to this reserve at every segment.

Proposition 1 (segment and terminal-brake admissibility). Suppose the desired path starts and ends at zero speed and satisfies the stated derivative bounds. Suppose it remains in $[\ell_{\min} + \delta_\ell, \ell_{\max} - \delta_\ell]$ and the following inequalities hold:

$$\begin{aligned} S_d + \tau_0 A_d &\le V_{\max}, \\ S_d + E_v &\le V_{\max}, \\ \left(1 + \frac{2\Delta\tau}{\tau_{\min}}\right) A_d + \frac{E_{v,0}}{\tau_{\min}} &\le A_{\max}. \end{aligned}$$

Then the modeled actuator satisfies its stroke, speed, command, and acceleration bounds throughout the segment and subsequent braking. The bounded-command brake brings the actuator to zero velocity in finite time.

Proof. The error convolution gives the stated bounds for the unsaturated path response. The speed, command, and acceleration inequalities ensure that this response also satisfies the bounded virtual model throughout the commanded segment. The length mismatch is bounded by $E_{\ell,0} + E_v H$, leaving at least $E_v^2/(2a_b)$ of stroke at the endpoint. Under braking, the deceleration magnitude is $\min\{A_{\max}, (b + |v_i^a|)/\tau_i\} \ge a_b$. Speed therefore decreases, the stated stopping-time and displacement bounds hold, and the remaining stroke reserve contains the entire braking motion. At the zero-speed event, the command switches to zero and preserves rest.

Corollary 1 (recursive actuator-constraint admissibility with a resting fallback). Define the resting set

$$\mathcal{S}_{\mathrm{rest}} = \{(\ell^a, v^a): \ell_{\min} \le \ell_i^a \le \ell_{\max}, \quad v_i^a = 0,\ i = 1, \dots, 4\}.$$

At each supervisor update from this set, a common four-actuator path is executed if Proposition 1's conditions hold for every actuator. Otherwise, the supervisor holds the current lengths with zero command. Each accepted segment followed by braking returns to $\mathcal{S}_{\mathrm{rest}}$, and holding preserves this set. By induction, an

admissible actuator action exists at every supervisor update for any requested target. This invariant-set statement guarantees an admissible actuator action, including a zero-progress hold. Tracking behavior and rigid-pose compatibility are evaluated separately in Sections 2.7 and 4.

The construction accommodates arbitrary received pose-request sequences and gains in $[0,1]$. Certification and the resting fallback preserve actuator admissibility independently of the information used to generate each request. Tracking performance under stale information is evaluated separately from this admissibility property.

2.6 Data-age modulation and gain control

Each patient-state packet records its sampling time t_s and arrival time t_a. At controller time t_c, only packets satisfying $t_a \le t_c$ are eligible. The newest arrived sample is either held or extrapolated using the two newest arrived snapshots. Under both information conditions, the consumed state retains its source sample time.

The Curvature variant first computes and rechecks a common certified path fraction α_c. The age-aware policy uses the governor decision time t_g and the source sample time t_s^* of the state most recently consumed by the adaptive controller:

$$\begin{aligned} a_g &= t_g - t_s^*, \\ \gamma(a_g) &= \max\{0, 1 - a_g/T_a\}, \\ \alpha &= \gamma(a_g)\alpha_c, \quad T_a = 0.250 \text{ s}. \end{aligned}$$

When no state is available, the policy sets $\gamma = 0$; the fixed-gain comparator uses $\gamma = 0.5$. Both timestamps, the gain, and the certified and applied path fractions are logged for each segment. Every applied common fraction is rechecked for all four actuators, with zero progress selecting the resting hold. Patient-state age thus regulates reference motion within the same actuator-admissibility checks. Tracking is compared with both the unmodulated Curvature governor and fixed attenuation.

2.7 Pose reconstruction and numerical realization

The four integrated actuator lengths need not lie exactly on the three-dimensional rigid-pose manifold. At each reconstruction, let $\mathbf{a}_i$ be the moving attachment in platform-local coordinates, $\mathbf{p}_0, \mathbf{Q}_0$ the segment-start pose, and $\mathbf{c}$ the fixed rotation center. Let $\boldsymbol{\xi} = (\eta_h, \eta_\phi, \eta_\psi)$ be the local pose increment from this segment reference: η_h is translation along the segment-start local vertical and the two angles parameterize the incremental world-axis rotations in the order below. These are not the global requested coordinates $\mathbf{x}$ in Section 2.1. The implemented map is

$$\begin{aligned} \mathbf{R}_{\text{inc}}(\boldsymbol{\xi}) &= \mathbf{R}_z(\eta_\psi)\mathbf{R}_x(-\eta_\phi), \\ \mathbf{p}(\boldsymbol{\xi}) &= \mathbf{c} + \mathbf{R}_{\text{inc}}(\boldsymbol{\xi})(\mathbf{p}_0 - \mathbf{c}) + \mathbf{Q}_0\mathbf{e}_y\eta_h, \\ f_i(\boldsymbol{\xi}) &= \| \mathbf{p}(\boldsymbol{\xi}) + \mathbf{R}_{\text{inc}}(\boldsymbol{\xi})\mathbf{Q}_0\mathbf{a}_i - \mathbf{b}_i \|, \\ \hat{\boldsymbol{\xi}} &\approx \operatorname{argmin}_{\boldsymbol{\xi}} 1/2 \sum_{i=1}^{4} [\, f_i(\boldsymbol{\xi}) - \ell_i^a]^2. \end{aligned} \tag{10}$$

The solver performs local damped least squares from the previous reconstructed pose in the current segment, with a zero relative pose at a new segment. Warm starts initialize a local assembly-branch search. For residual vector $\mathbf{e} = \mathbf{f} - \ell^a$, each trial solves $(\mathbf{J}^T\mathbf{J} + \lambda\mathbf{I})\Delta\boldsymbol{\xi} = -\mathbf{J}^T\mathbf{e}$. Forward differences use 10^{-5} m in heave and 10^{-4} rad in each angle. The damping starts at 10^{-7}, decreases by a factor of four after an objective reduction (floor 10^{-12}), and increases tenfold after a rejected trial. The production call allows thirty iterations. Immediate termination requires $\| \mathbf{e} \|_\infty \le 1$ micrometre; final acceptance requires $\| \mathbf{e} \|_\infty \le 0.5$

mm. Both this maximum norm and $\| \mathbf{e} \|_2/2$ are returned by the solver. Exact pseudocode and incompatibility controls are in Online Resource 2, Section S5.

Pose compatibility is an execution-acceptance condition. An out-of-tolerance solve increments the incompatibility counter while the last iterate remains displayed; that execution is excluded. All eighty accepted Unity executions satisfy this condition, comprising the original sixty-run matrix and the targeted twenty-run matrix in Section 3.4. The maximum discrepancy between integrated lengths and reconstructed attachment distances supplies the next segment's initial-error bound, connecting numerical reconstruction to the subsequent actuator certificate.

The native controller and patient schedules are retained. Each render-frame duration is integrated with fourth-order Runge–Kutta substeps of at most 2 ms, including remaining time at segment boundaries. Braking accounts for saturation transitions and the exact zero-speed event. Commands and state updates are checked before numerical clamping, and limit events invalidate an execution.

2.8 Generative AI use

An OpenAI Codex assistant supported mathematical drafting, implementation changes, experiment orchestration, analysis scripts, figure preparation, and manuscript drafting. The workflow also used procedural guidance from a scientific-agent skills library [9]. Numerical outputs were generated by Unity and the specified numerical solvers. The authors retain responsibility for validating the equations, implementation, analysis, references, and final manuscript.

3 Experimental design

3.1 Shared-input benchmark with matched geometric baselines

The shared-input bank contains 540 coordinate-wise paths generated with .NET Random seed 20261028: 180 each for heave, pitch, and roll. Start and goal coordinates are sampled independently and uniformly. Heave starts lie in [−0.010,0.010] m and goals in [−0.018,0.018] m; angular starts lie in [−0.035,0.035] rad and goals in [−0.055,0.055] rad. The other coordinates are fixed at zero, and scene geometry determines the four attachment paths. Every generated path enters the comparison, independently of certificate success, admitted progress, or tracking outcome.

Each method receives the same four-joint path, reconstructed geometry, prescribed initial length offset (0, 0.1, or 0.5 mm), and zero initial velocity. Every complete decision measures the float-rounded initial discrepancy, computes the response/braking reserve, searches four actuator prefixes, selects their minimum, and rechecks the common prefix. It retains the same resting fallback. The comparison includes four analytic methods and ten search methods, giving fourteen methods in total. The ten search methods use norm-envelope or projected-envelope DirectSearch at five subinterval counts. Conservative and Curvature retain their whole-vector bounds; AxisConservative and AxisCurvature use the projected-radius bounds in Section 2.2. The original sixty closed-loop executions use norm-envelope methods. The targeted matrix in Section 3.4 adds AxisConservative to the closed loop; AxisCurvature and the search-resolution scans remain offline comparisons.

Calculations run synchronously in Unity 2022.3.62f2c1 Editor Edit mode with `-batchmode -nographics -executeMethod` on the stated Intel Core i7-14650HX/Windows 11 workstation. This synchronous calculation benchmark runs outside the Play-mode frame loop. Each path–offset–method combination receives two warmups and twenty timed repeats; method order rotates with path and repeat. The unified batch contains 453,600 timed decisions. The timed block includes field extraction, budgets/reserves, four searches, and the common-prefix recheck; it excludes CSV parsing/formatting, actuator integration, pose

fitting, patient updates, and rendering. All within-batch elapsed times, including maxima, are retained. These elapsed-time distributions characterize the specified implementation and host.

Each of the 540 paths contributes once to summaries of certified progress, the dimensionless path fraction returned after the common-prefix recheck. Its twenty identical progress outputs are timing repeats. Gains use ratios of method means, and Figure 1 reports paired path differences. The contribution attributed to axis projection uses the additive progress increment, $(\overline{\alpha}_{\mathrm{AxisC}} - \overline{\alpha}_{\mathrm{C}})/(\overline{\alpha}_{\mathrm{K}} - \overline{\alpha}_{\mathrm{C}})$, with all means taken on the same paths. Table 2 reports selected 0.1-mm methods, Table 3 gives coordinate strata, and Online Resource 2, Tables S11 and S16, contains all methods and offsets. Online Resource 3 supplies all path vectors, actual virtual geometry coordinates, supervisor source/entry points, golden outputs, timing records, and an independent asset-free replay. Online Resource 2, Section S14, specifies the pinned Mono/compiler configuration for that replay. The separate core-certificate experiment in Online Resource 1, Section S2, excludes the complete initial-error/stopping supervisor.

3.2 Initial-error and braking checks

The numerical bank crosses twelve specified synthetic paths with three certificates and signed initial position and velocity discrepancies. Their magnitudes are 0.20 mm and 0.01 mm/s, respectively, with both signs tested. Four response times (0.030, 0.040, 0.060, and 0.090 s) and two integration steps (2 ms and 1 ms) complete the design. The resulting 1,152 cases include both endpoints of the response-uncertainty interval. Each case checks positive certified progress, length, speed, path-phase acceleration, and subsequent stopping displacement. For Curvature, 12 paths × 2 initial-position-error signs × 2 initial-velocity-error signs give 48 certified cases. Rechecking each at the five multipliers $\{0,0.25,0.5,0.75,1\}$ gives 240 modulation checks; these are certificate checks before response-time/integration-step expansion. In 96 cases, the event-aware brake is compared independently with fine-step numerical integration across both velocity signs, saturation transitions, and three integration durations. Detailed geometry, tolerances, and numerical outputs are provided in the supplementary information.

The boundary bank tests both stroke limits, nominal initial headrooms of 0.2, 0.408, 0.5, 0.75, 1.2, and 3 mm, outward initial errors of 0, 0.2, and 0.5 mm, three certificate methods, and response constants of 0.03, 0.06, and 0.09 s. Initially in-stroke states form the evaluation domain. Complete, initial-error-omitting, and braking-term-omitting checks receive the same synthetic scalar length paths, with 1-ms RK4 response integration and event-aware stopping without stroke clipping. Independent stopping controls begin with 0, half, or all of the stopping-distance bound at both limits. Together, these controls identify when each reserve changes progress, fallback selection, and the final stopping margin. Online Resource 2, Sections S6–S7, supplies endpoints, terminal velocities, minimum margins, held cases, and the complete timing bank.

3.3 Original matched closed-loop matrices

The Unity matrix uses three patient profiles in five model–seed conditions: ModerateImpairment with seeds 20261031–20261033, MildImpairment with seed 20261034, and Healthy with seed 20261035. All conditions use Standard terrain, ACOR, the ModelQuaternion angular source, and RadianReference normalization, without additional measurement-noise injection. Inherited profile-dependent perturbations and center-of-mass (COM) drift remain part of the digital model. Patient-update eligibility is configured at 25 Hz, while actual sampling and consumption follow the native schedules and are logged.

D80 compares Conservative, Curvature, and DirectSearch under held and capped-extrapolated feedback, yielding thirty executions. D240 compares Curvature, age-modulated Curvature, and Curvature with a constant 0.5 gain under the same feedback conditions, yielding another thirty executions. The scheduled packet-delivery delay ranges are 0–80 ms and 0–240 ms, respectively. Delivery delay depends on packet number and seed phase; corresponding packet indices receive identical scheduled delays. Actual

consumption ages and terminal packet counts are retained as separate quantities. Execution order is shuffled within model–seed and delay blocks using randomization seed 20261029.

Each execution comprises a 15-s adaptation phase and a 90-s recording phase containing 4,500 primary fixed-step samples. Tracking root-mean-square error (RMSE) is computed against terrain inputs over the primary trace. State age, state error, segment-level gain, and progress are averaged over the same first-to-last acquisition timestamp window. Whole-run timestamp, certificate, and limit-event checks also cover adaptation and return. Acceptance requires completion, zero runtime errors, and zero pose-fit acceptance failures. Path command, speed, acceleration, and stroke-limit events must remain absent beyond the stated tolerances. The configuration matrix and seven key implementation source snapshots were frozen before collection; configuration and run identifiers are retained.

Each execution is the reporting unit. Five model–seed conditions crossed with two feedback modes provide ten paired contrasts per comparison, and the complete execution-level results and matched changes are retained. The age policy and fixed 0.5 gain define distinct allocation rules, with their mean gains reported explicitly. The shared-input benchmark measures certificate computation under common inputs; the closed-loop matrices quantify total execution behavior with evolving reference trajectories.

Time-history illustrations use ModerateImpairment, seed 20261031, with held feedback, the first condition in the frozen matrix. This illustration rule was selected after collection, and the execution-level summaries retain all sixty runs. Each tracking error uses its execution's logged terrain input and native timestamps.

Common-reference re-scoring quantifies sensitivity to the error-reference choice. Stored D80 outputs are scored against the Conservative run's terrain input on the common 4,500-sample acquisition grid while retaining their original trajectories. Reference RMS differences bound the possible RMSE change by the triangle inequality. Native frame intervals, decision counts, and observed rest durations characterize the accompanying input and scheduling differences. This analysis therefore estimates the sensitivity of configuration-matched total execution behavior.

Per-coordinate active intervals use an absolute terrain request above 1° as a numerical selection threshold. Active-interval RMSE, actual angular total variation, peak response, coordinate-selection shares, and time-weighted observed rest (logged maximum virtual speed at most 0.001 mm/s) complement full-window errors. Rest duration uses native frame observations. Online Resource 2, Sections S8–S9 and Table S13, provides these diagnostics for all sixty executions.

3.4 Axis-aware closed-loop control and decision budgets

A targeted D80 matrix crosses Curvature and AxisConservative with the same five model–seed conditions and two feedback modes, giving twenty accepted executions. The two methods use identical virtual response dynamics, initial-error and stopping reserves, 0.8-s segments, Standard terrain, and native schedules. The added axis mode replaces the attachment radius throughout the Conservative certificate by its rotation-axis projection. Both methods log the complete four-actuator geometry, integrated initial states, nominated target, and selected coordinate after each decision. The run order was fixed before collection using seed 20261111. Acceptance retains the original 4,500-sample and numerical-check requirements; all attempts and exclusion reasons are retained in Online Resource 2, Section S16.

The closed-loop contrasts estimate configuration-matched total execution effects. A separate same-state replay evaluates Conservative, AxisConservative, Curvature and 16-subinterval DirectSearch on every captured decision input with identical complete reserves. Runtime progress is verified against the captured output at absolute tolerance 2e−5. For Curvature decisions, four diagnostic bisections separately restrict desired acceleration, command speed, response-speed enclosure or reserved stroke, retaining positive geometric denominators. Their individual caps identify the binding channel; a cap computed after relaxing

other channels is not an executable command. Additional reserve-omission diagnostics remove one length-reserve term at a time while retaining the same response and acceleration budgets.

The representative decision is selected by the largest absolute nominated pitch target in the first model–seed, held-feedback Curvature execution, with earliest time breaking a tie. This descriptive selection is applied after collection. Desired derivatives are sampled at 1,001 normalized times on its four recorded joint paths to compare the analytic enclosure with the evaluated geometry. The full comparison and diagnostic definitions are given in Sections S16–S17.

Online Resource 1. Supplementary numerical verification, analytical boundary-reserve examples, experimental configurations, all sixty execution-level results, and the retained-file data dictionary (ESM_1.pdf).

Online Resource 2. Pose reconstruction, boundary responses and stopping, shared-input timing, input/schedule sensitivity, activity-conditioned results, inherited-model details, and the targeted axis-aware comparison with decision budgets (ESM_2.pdf).

Online Resource 3. Asset-free reproduction of the complete 540-path supervisor benchmark, including geometry, inputs, parameters, source, raw timing records, golden progress, analysis, and integrity checks (ESM_3_Path_Benchmark.zip).

Data and code availability: The complete shared-input path benchmark and author-developed replay source accompany this article as Online Resource 3. The full Unity closed-loop datasets, configurations, and source snapshots, including the targeted twenty-run matrix and captured decision inputs, are available from the corresponding author on reasonable request. Third-party assets retain their original licenses and are not required by the path-benchmark replay.

4 Results

4.1 Geometric allowance and complete-supervisor decision time

Curvature admits 7.69% more mean certified progress than AxisConservative, the inexpensive axis-aware norm baseline at a 0.1-mm initial length offset (Table 2): 0.349797 versus 0.324806. Their median decision times are 53.50 and 23.00 μs, respectively, quantifying the computation exchanged for this additional allowance. The progress gain remains 7.69%–7.71% across the 0–0.5-mm offsets. Whole-vector Conservative admits mean certified progress 0.250049, giving a 39.89% difference whose geometric sources are separated below. Every one of the 453,600 decisions passed the common-prefix recheck, and all interior-bank decisions admitted positive progress. Section 4.2 resolves shortening and resting fallback near the stroke boundaries.

Table 2. Complete-supervisor comparison at a prescribed 0.1-mm initial length offset. Each row contains 540 shared paths and 10,800 timed calls with the same initial-error, response and terminal-braking reserve. Axis modes use projected-radius bounds; search suffixes specify subinterval count. Pooled-call times cover the decision block in Section 3.1. These are offline comparisons; the original sixty closed-loop executions use the norm-envelope methods.

Method	Mean progress	Zero (%)	Median (µs)	P95 (µs)	Maximum (µs)
Conservative	0.250049	0.0	44.40	72.50	2833.2
AxisConservative	0.324806	0.0	23.00	36.30	1236.0
Curvature	0.349797	0.0	53.50	82.70	2375.8
AxisCurvature	0.356402	0.0	33.30	51.30	1097.3
Axis search (1)	0.357757	0.0	150.80	237.91	2434.7
Axis search (16)	0.359974	0.0	1308.10	1858.04	6199.2

Axis projection alone recovers 74.95% of the additive Curvature-over-Conservative progress increment at 0.1 mm. For Roll, its share is 91.94%. The large whole-vector comparison therefore combines removal of a loose radius envelope with directional curvature refinement. Table 3 separates these effects: Curvature's gains over AxisConservative are 11.02%, 9.72%, and 4.28% for Heave, Pitch, and Roll. Across these strata it improves 167, 171, and 146 paired paths, with 13, 9, and 34 ties and no negative difference (Figure 1b). The projected Curvature variant additionally tightens its remainder and achieves 9.73% pooled gain over AxisConservative; its result belongs to the projected offline variant.

Table 3. Coordinate-wise progress at 0.1 mm, with 180 shared paths per coordinate. Axis norm denotes AxisConservative; both gain columns use its mean as denominator. Curvature gains against the whole-vector baseline are 11.02%, 15.39%, and 103.50%, respectively; they are not the residual curvature increment. All progress values are dimensionless; timing repeats do not increase the path count.

Coordinate	Conservative	Axis norm	Curvature	Curvature gain (%)	AxisCurvature gain (%)
Heave	0.315372	0.315372	0.350111	11.02	11.02
Pitch	0.210517	0.221393	0.242909	9.72	9.84
Roll	0.224257	0.437653	0.456372	4.28	8.74

Two geometric refinements explain the progress gain. The whole-vector radius includes attachment components parallel to the rotation axis; removing them explains much of the Roll increment (Figure 1c). AxisConservative still bounds every length derivative by the attachment-motion magnitude. Curvature instead uses the initial radial directional slope and bounds its change with a curvature remainder. AxisCurvature also removes the parallel component from that remainder. The paired paths therefore distinguish the allowance recovered by axis projection from the additional allowance recovered by directional curvature.

Search resolution materially changes the decision-time comparison (Figure 1d). DirectSearch with whole-vector bounds and one subinterval admits mean progress 0.353098, versus 0.357713 with sixteen; its median time is 148.60 rather than 1304.70 µs. Relative to Curvature with whole-vector bounds, the measured ratios are 2.78 and 24.39, respectively. With projected-radius bounds, one- and sixteen-subinterval AxisDirectSearch variants admit 0.357757 and 0.359974; their median times are 4.53 and 39.28 times AxisCurvature's. These implementation-specific ratios locate each resolution on the measured certified-progress–decision-time trade-off. Table S16 reports all five counts under both envelopes, including P95 and maxima.

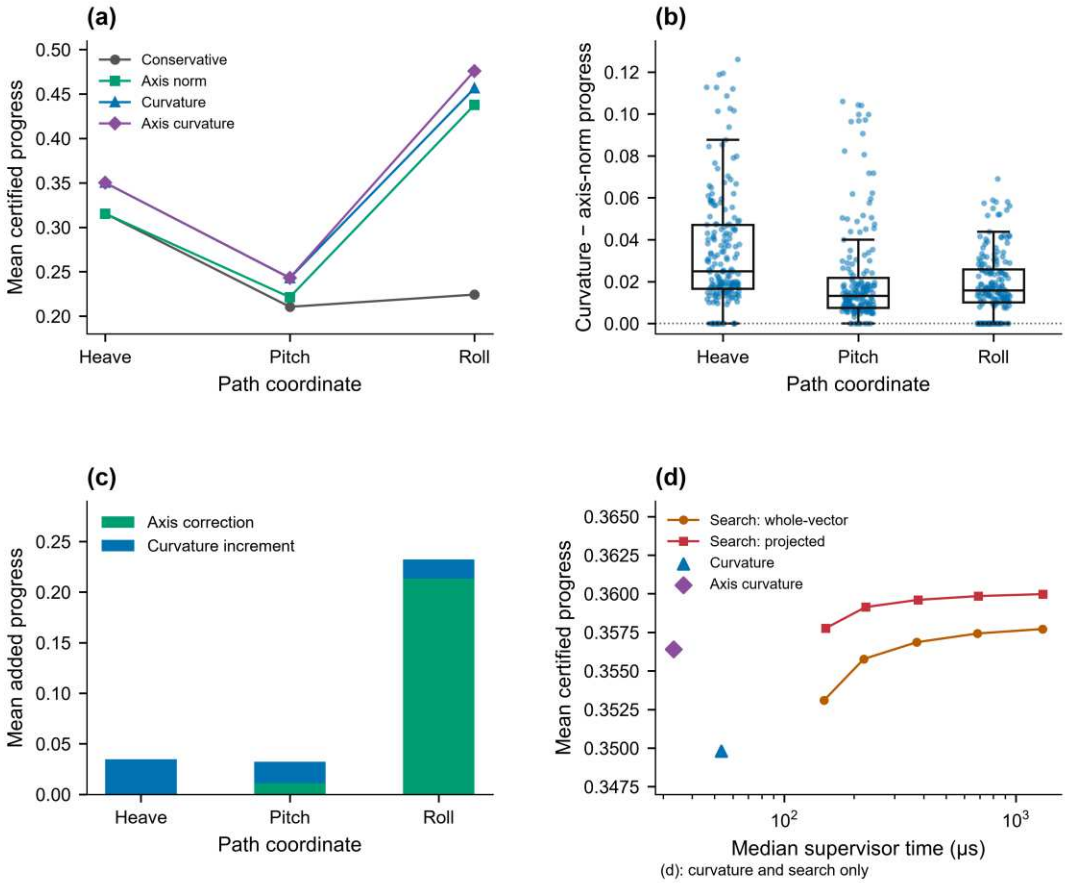

Fig. 1 Separating geometric baseline tightening from curvature refinement and search cost at the 0.1-mm prescribed initial offset. (a) Coordinate-wise mean progress for four analytic modes. (b) Paired Curvature minus AxisConservative progress, 180 paths per coordinate; boxes show median and interquartile range, whiskers extend to 1.5 interquartile ranges, and points retain all paths. (c) Mean additive progress increments from whole-vector Conservative to AxisConservative and from AxisConservative to Curvature. (d) Mean progress versus pooled median decision time for both search envelopes at 1, 2, 4, 8, and 16 subintervals and only the two curvature variants; time uses a logarithmic axis. Panel (d) has its own legend and a restricted progress range; Conservative and AxisConservative are compared in panel (a) and Table 2, not plotted in panel (d). All timing values come from the same completed batch. AxisCurvature and the search scans are offline variants

4.2 Initial-state and finite-stop numerical checks

All 1,152 initial-error response cases passed the specified checks with positive returned progress. Observed lengths, including braking endpoints, ranged from 0.464265 to 0.520866 m. Maximum path-response speed was 0.003361 m/s, and maximum path-phase acceleration was 0.033375 m/s^2, both within the assumed limits. All 240 Curvature progress-modulation checks passed the specified numerical acceptance criteria.

All 96 braking comparisons against independent fine-step integration passed the specified checks. Maximum velocity and displacement discrepancies were 3.95×10^{-9} m/s and 1.35×10^{-9} m, respectively. These comparisons verify the numerical response and event-handling implementation; Proposition 1 and Corollary 1 provide the analytical admissibility argument.

The boundary bank yielded 288 complete-reserve response cases, comprising 180 positive-progress responses and 108 stationary fallbacks, with no negative stroke margin beyond numerical tolerance. Omitting

the initial-error term caused 54 violations among the corresponding responses. With an outward 0.2-mm initial error and 0.75-mm nominal headroom, Curvature shortened the motion to approximately 0.07915 at both limits. At the upper limit its full-stop endpoint was 629.58967 mm with a 0.41033-mm minimum margin; at the lower limit these values were 380.41029 mm and 0.41030 mm. Reducing nominal headroom to 0.5 mm selected a stationary fallback, with actual lengths 629.69995 and 380.29999 mm, respectively.

The reserve controls separate two mechanisms. Initial displacement error and residual stopping displacement act outward in the response bank; removing the initial-error term exposes the 54 violations reported above. The braking-omission responses remain within stroke in that bank, while independent terminal-state controls isolate the stopping requirement. At 0.5 mm/s terminal speed and a 0.09-s response constant, stopping from either stroke limit crosses it by approximately 6.16 μm. Retaining the 7.5-μm braking bound leaves approximately 1.3 μm of observed margin. Figure 2 distinguishes progress decisions, initial-error controls, and stopping controls; Online Resource 2, Section S6, provides full endpoints and minimum margins.

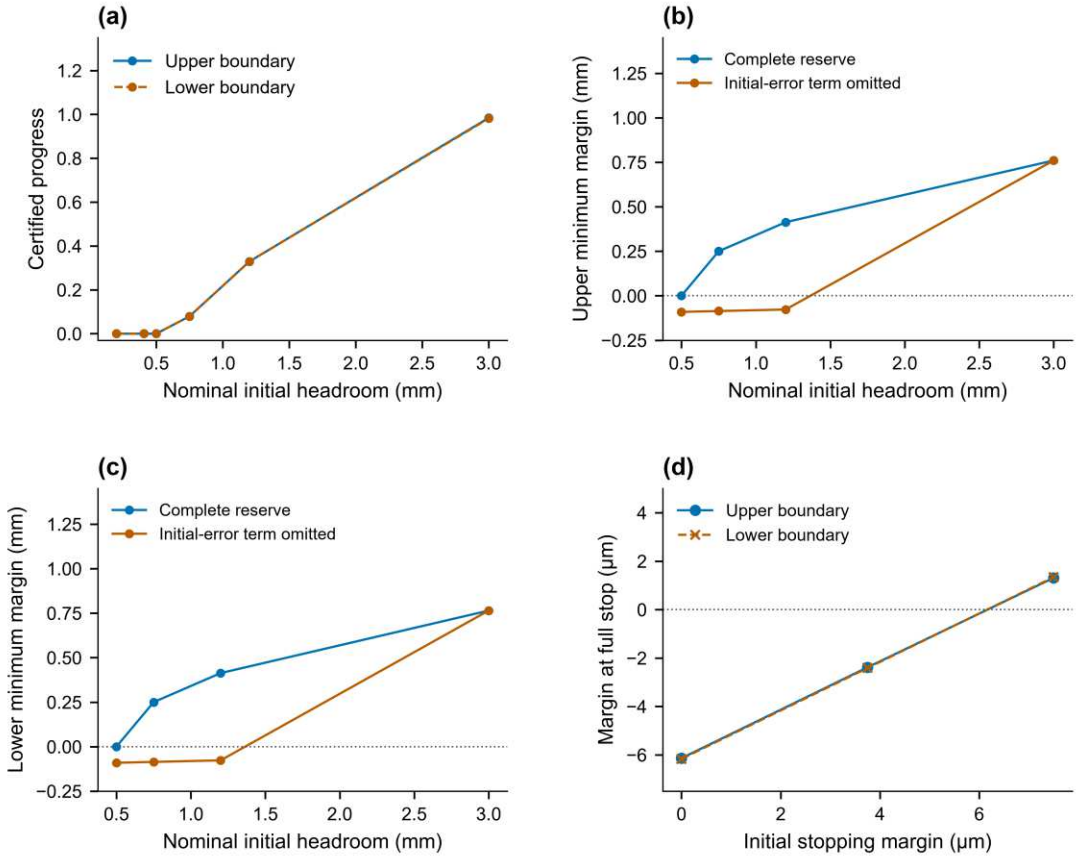

Fig. 2 Boundary-focused virtual actuator controls. (a) Curvature progress versus nominal initial headroom at both limits, with a 0.2-mm outward initial error and 0.09-s response time; zero progress denotes stationary fallback. (b, c) Minimum observed upper/lower stroke margin through response and stopping, with a 0.5-mm outward initial error, comparing complete reserves with omission of the initial-error term. (d) Full-stop margins at 0.5 mm/s outward speed and 0.09-s response time, starting with zero, half, or all of the 7.5-μm stopping-distance bound. Points are computed cases and lines connect them for reading. These synthetic scalar controls use no stroke clipping and are separate from the Unity patient-model executions

4.3 Continuous virtual actuator response and constraint margins

In the original matrix, all sixty executions completed, each with 4,500 primary observations in the recording phase. No runtime errors, pose-fit acceptance failures, or path command, speed, acceleration, or stroke-limit

events were recorded. Across complete retained runs, maximum actual actuator speed was 0.0033603 m/s, and maximum actual acceleration was 0.044974 m/s^2. The largest pose-fit residual was 0.12067 mm, and the largest measured segment-start length discrepancy was 0.10142 mm. Every segment began from a resting actuator state, as required by the supervisor.

All fifty packet-delay pairing checks passed, with no mismatches at shared packet indices. The analysis retained fifty method/policy contrasts and thirty held-versus-extrapolated feedback contrasts. The largest decision-age reconstruction discrepancy was below 2.7×10^{-14} s; gain and applied-fraction discrepancies were below 4.6×10^{-8}. At corresponding acquisition indices, maximum terrain-input differences were 0.06753° for D80 and 0.06364° for D240. These comparisons matched model–seed and packet-delivery configurations while retaining native software schedules.

Table 4. Whole-run actuator-state and reconstruction checks. Extrema come from retained per-run summaries and include adaptation, recording, and return. The acceleration entry is a summary maximum, not a reconstructed time-history derivative. All limits are declared simulation assumptions or numerical acceptance thresholds.

Quantity	Declared simulation limit or acceptance threshold	Observed extrema over sixty complete runs
Integrated virtual actuator length (mm)	380–630	443.049–494.953
Absolute virtual actuator speed (mm/s)	9.5	3.36030 maximum
Absolute virtual actuator acceleration (mm/s²)	50	44.97352 maximum
Platform-pose fit residual (mm)	0.5	0.12067 maximum

The recorded full virtual-actuator update, including integration and pose fitting, had D80 per-frame medians of 137.3–141.1 µs across the three methods. P95 values were 198.4–228.9 µs and observed maxima were 3,841.4–9,246.8 µs. These native-frame timings describe complete updates on each execution's own input, complementing the shared-input decision timings in Section 4.1. Online Resource 2, Section S12 and Table S14, defines the measured block and reports every delay/method group.

4.4 Tracking against the whole-vector baseline

Figure 3 compares requested and realized platform angles and their tracking errors for the selected D80 condition. Table 5 gives the corresponding absolute RMSE and applied progress. Curvature increased mean applied progress from 0.08636 to 0.16907 and reduced roll RMSE from 1.78899° to 1.66837°, while pitch RMSE was similar (4.61391° and 4.61422°). DirectSearch yielded roll and pitch RMSEs of 1.58627° and 4.59101° in this condition. The time histories show when the methods diverge; Figure 4 places the illustration within all ten matched contrasts.

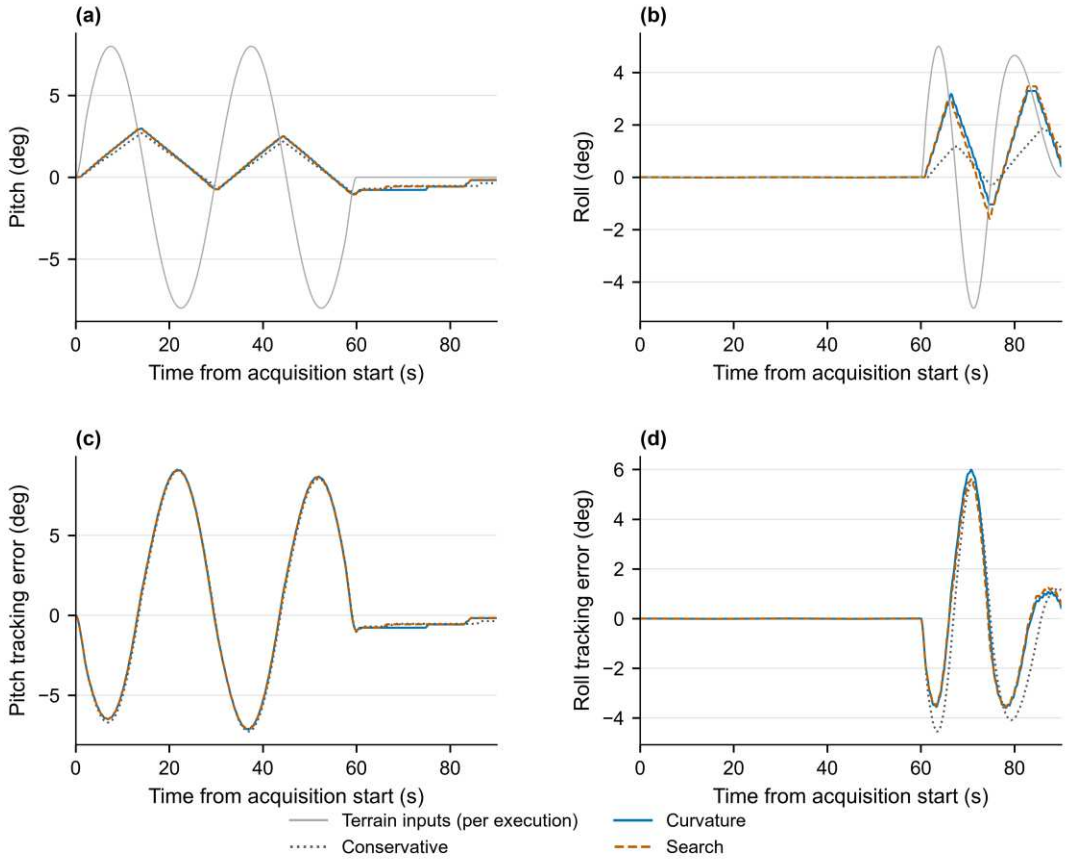

Fig. 3 D80 tracking time histories for ModerateImpairment, seed 20261031, held feedback. (a, b) Realized pitch and roll for the three certificate methods, with each execution's terrain input in gray. (c, d) Realized-minus-requested pitch and roll errors, computed against each method's own logged terrain input. Native timestamps are aligned to the start of each recording window; no smoothing or interpolation is applied. Closed-loop terrain inputs can differ between executions

Table 5. Absolute tracking and motion-allocation results for the D80 condition illustrated in Figure 3. RMSE uses all 4,500 primary samples per execution; segment progress is averaged within the recording window. Complete execution-level results are retained in Online Resource 1.

Method	Pitch RMSE (deg)	Roll RMSE (deg)	Mean applied progress
Conservative	4.61391	1.78899	0.08636
Curvature	4.61422	1.66837	0.16907
DirectSearch	4.59101	1.58627	0.17167

Under D80, Curvature reduced roll-tracking RMSE by 6.52%–13.78% relative to Conservative across the ten configuration-matched comparisons in Figure 4. Pitch-RMSE changes ranged from −2.47% to +0.007% in the same comparisons. Mean applied segment fractions were 0.1667–0.1877 for Curvature and 0.0862–0.0979 for Conservative. These configuration-matched executions associate greater motion allocation with lower roll-tracking error; they estimate total behavior of the coupled platform, terrain input, and native schedules. The comparison is against whole-vector Conservative; the incremental closed-loop comparison against AxisConservative is given in Section 4.6.

DirectSearch reduced roll RMSE by 9.83%–13.21% and pitch RMSE by 0.47%–2.61% relative to Conservative. Its tracking order against Curvature depends on the condition. Together with the shared-input timing results in Section 4.1, these outcomes characterize the available tracking–computation choices; closed-loop trajectories and common-input certificate decisions remain separate measurement objects.

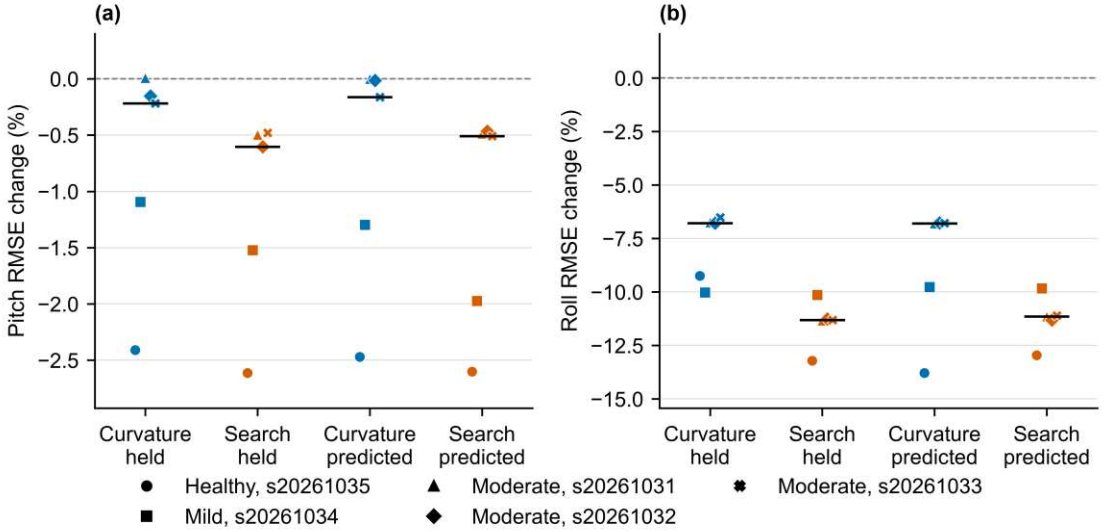

Fig. 4 D80 configuration-matched tracking changes relative to Conservative. Each marker represents one of five model–seed contrasts under held or extrapolated feedback; black strokes show descriptive medians. Negative changes indicate lower RMSE. Pitch and roll are shown separately, and all ten contrasts per certificate are retained

The roll-tracking result persists under common-reference re-scoring: all ten Curvature contrasts retain lower roll RMSE, with reductions of 6.59%–13.71%. The largest D80 pointwise terrain-input difference is 0.06753°, and one near-tied pitch ordering changes under re-scoring. Across the twenty alternative-versus-Conservative contrasts, mean native frame-interval differences range from −0.10079 to +0.04625 ms. This analysis establishes robustness of the observed roll ordering to the score-reference choice while retaining native scheduling as part of the total execution comparison.

Absolute response and motion exposure identify the operating regime. For requests peaking at 8° pitch and 5° roll, Curvature reaches 2.976° pitch, or 37.2% of the requested peak; pitch RMSE is 4.614° against a request RMS of 4.612°. Pitch receives 70% of decisions and roll 30%, with 0.333 s of time-weighted observed rest. Table 6 adds active-interval error and realized angular travel to full-window RMSE. These measurements describe constrained partial terrain presentation and the motion allocated by the one-coordinate supervisor.

Table 6. Activity-conditioned performance for the condition in Figure 3. Active intervals use per-coordinate requests above 1° (54.46–54.52 s for pitch and 24.40–24.42 s for roll across methods). Travel is total variation of realized angle over the complete 90-s trace. Observed rest uses native-frame time weights and the 0.001-mm/s virtual-speed threshold, not a count of exact internal hold events. All sixty execution-level values are supplied in Online Resource 2.

Method	Active pitch RMSE (deg)	Active roll RMSE (deg)	Pitch travel (deg)	Roll travel (deg)	Observed rest (s)
Conservative	5.890	3.400	12.786	5.793	0.300
Curvature	5.882	3.132	14.403	14.828	0.333
DirectSearch	5.859	2.972	14.480	15.535	0.278

4.5 Information-age allocation and state extrapolation

Figure 5 links the delivered information to the applied motion allocation in the selected D240 age-policy execution. Decision age differs from scheduled packet delay because packet consumption follows the native controller schedule. Larger decision ages reduce the logged policy gain, which scales the certified progress before application. The fixed-gain line is the configured 0.5 reference; cross-execution comparisons are reported in Figure 6.

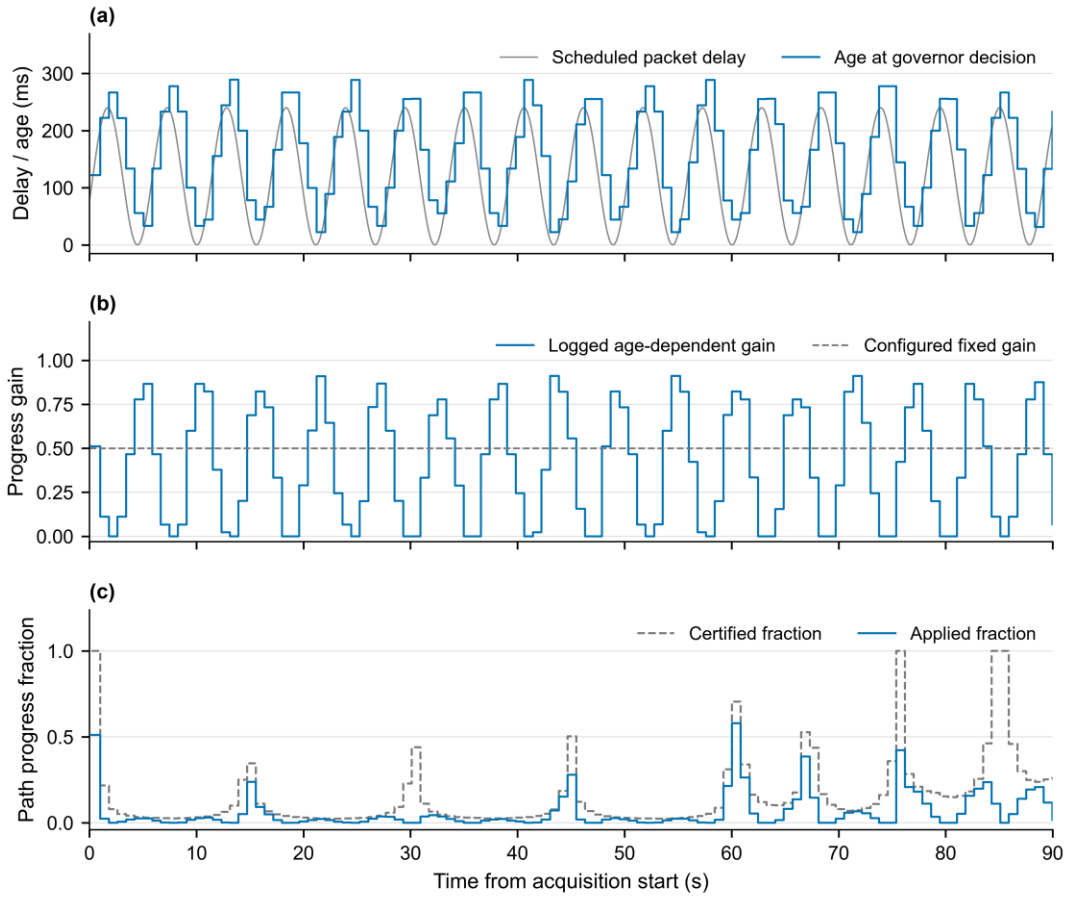

Fig. 5 Timestamp-to-progress mechanism in D240 age-modulated Curvature, ModerateImpairment, seed 20261031, held feedback. (a) Scheduled packet-delivery delay at packet sampling times and consumed-state age at governor decisions. (b) Logged age-dependent gain, with the configured fixed gain of 0.5 as a reference rather than a separate closed-loop trajectory. (c) Certified and applied path-progress fractions at segment decisions. Decision signals are shown as native step traces; no smoothing is applied

Figure 6 shows how timestamp-derived gain modulation reallocates motion under D240. The age policy's mean gain ranged from 0.3841 to 0.4012, with mean applied fractions of 0.0595–0.0655. Corresponding fraction ranges were 0.1651–0.1875 for unmodulated Curvature and 0.0793–0.0884 for fixed 0.5 attenuation. Relative to unmodulated Curvature, age modulation accompanied roll-RMSE increases of 5.58%–14.59% and pitch changes of −0.063% to +2.183%. These results quantify the tracking changes accompanying the different motion allocations.

Against fixed 0.5 attenuation, the age policy yields 0.49%–1.26% lower pitch RMSE across ten contrasts, with roll-RMSE changes of −1.38% to +1.77%. The two policies combine different allocation rules and different mean gains, so the comparison quantifies their total policy effect. Both retain the same actuator-admissibility checks.

Extrapolation improves the delivered state representation: mean consumed COM-state error decreases by 50.5%–70.9% across all thirty held-versus-extrapolated configuration pairs. The resulting pitch-RMSE changes span −0.443% to +0.644%, and roll changes span −4.994% to +0.613%. Thus, state-estimation accuracy and terrain tracking measure distinct responses of the coupled controller and actuator supervisor. Online Resource 1 retains every execution and matched contrast.

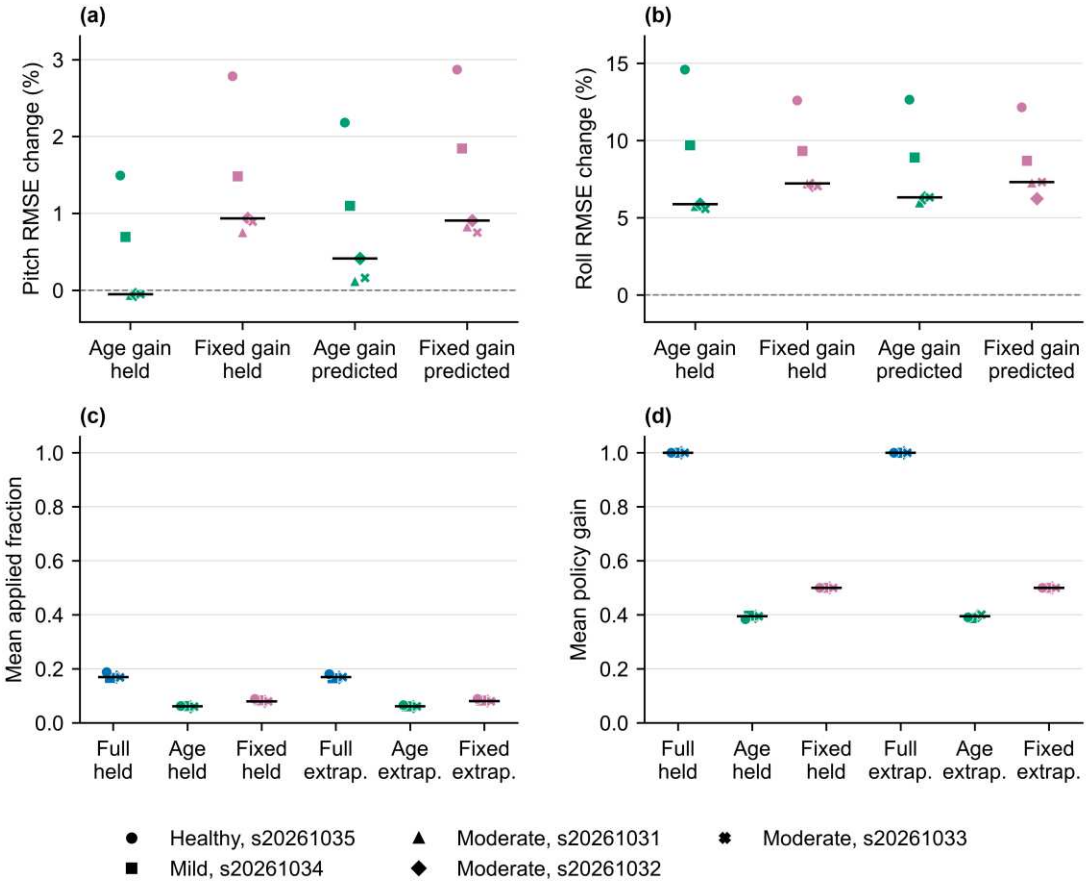

Fig. 6 D240 motion-allocation and tracking trade-off. Top panels show RMSE changes relative to unmodulated Curvature; bottom panels show each execution's mean applied fraction and policy gain. Full, Age, and Fixed denote unmodulated Curvature, age modulation, and fixed 0.5 attenuation; extrap. denotes extrapolated feedback. Markers identify the five model–seed settings, and black strokes are descriptive medians. Fixed attenuation is 0.5 and is not matched to each age-policy run's mean gain

4.6 Axis-aware comparison and the active motion budget

The targeted twenty-run matrix compares Curvature directly with AxisConservative under D80. Curvature's roll RMSE was 2.62%–3.70% lower than AxisConservative in all ten paired contrasts. Pitch-RMSE changes span -2.35% to -0.37%. Figure 7 retains all ten contrasts. Shared packet indices have identical scheduled delays, and re-scoring Curvature against each AxisConservative terrain trace gives roll-RMSE changes of -3.81% to -2.68%. These are configuration-matched execution effects under native scheduling. Replay on identical captured decision inputs measures certified progress without changing the path or initial state. Replayed and logged progress agree exactly in this batch (Online Resource 2, Section S17).

Across 1100 eligible primary Curvature decisions, 1052 shorten the requested path. In every shortened case, the desired-acceleration cap equals the fraction admitted by the full certificate. The robust desired-

acceleration budget is 0.016667 m/s², one third of the declared virtual-actuator acceleration limit. In the representative pitch decision, a 6.503° residual yields progress 0.028897, or a 0.188° accepted increment along the nominated pose path. The separate desired-acceleration, command-speed, response-speed and reserved-stroke caps are 0.028897, 0.062811, 0.070875 and 1.000000. Thus speed and stroke headroom coexist with an active acceleration certificate.

On this same decision, the 16-subinterval comparator admits 0.029084, 0.65% above Curvature. The sampled maximum desired acceleration is 0.016500 m/s², close to its analytic budget. The initial-error, response and braking length reserves are 0.0523, 0.4000 and 0.0075 mm; omitting each separately leaves the admitted progress unchanged, with a complete-certificate stroke margin of 73.13 mm. These reserves protect boundary cases rather than determine this interior decision. During the first 60-s pitch phase, all 74 eligible decisions select pitch. The representative attenuation therefore occurs under active pitch allocation, with the robust acceleration envelope and repeated zero-speed 0.8-s segments setting the local motion rate.

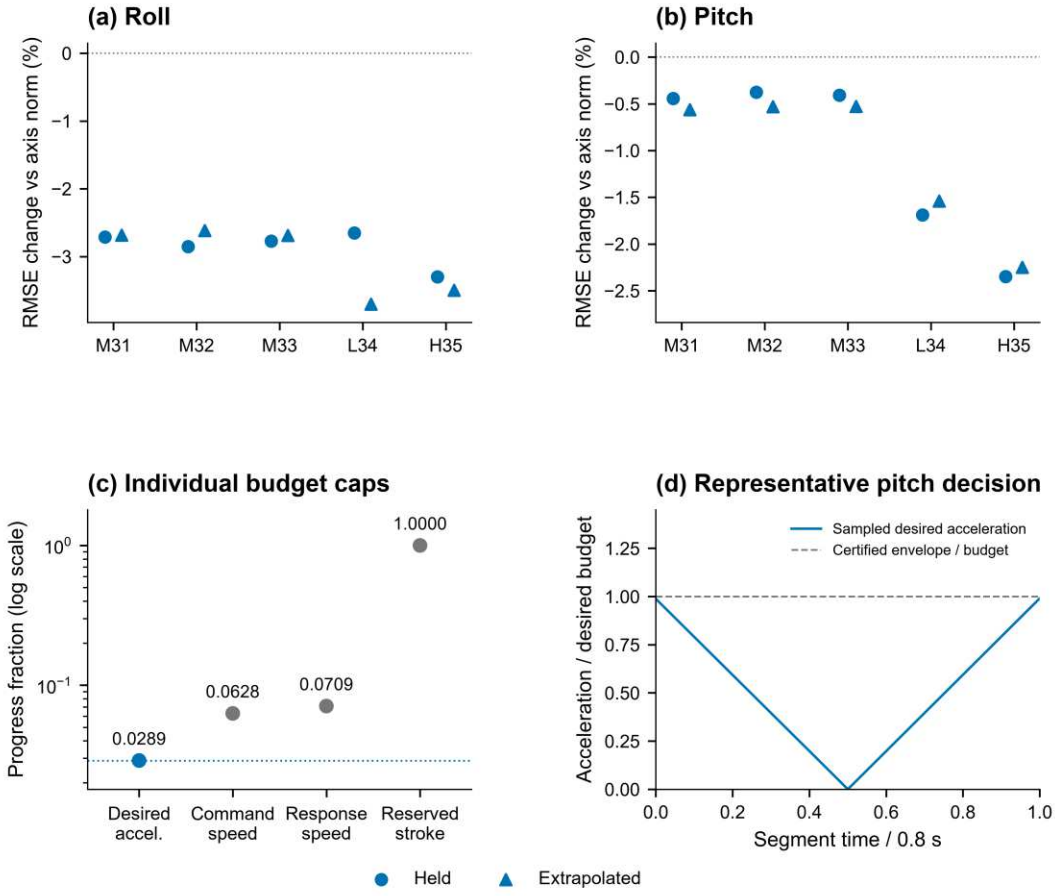

Fig. 7 Targeted axis-aware closed-loop comparison and a representative pitch-decision budget. (a, b) Curvature RMSE changes relative to AxisConservative in all ten configuration-matched contrasts; negative changes indicate lower error. Symbols identify held or extrapolated feedback in the five model–seed settings. (c) Progress caps obtained by retaining one certificate channel at a time on the same recorded four-actuator input; the dotted line is the full Curvature decision. These relaxed caps are diagnostics, not executable commands. (d) Maximum desired acceleration across four joint paths sampled at 1,001 segment times, normalized by the desired-acceleration budget, with the analytic enclosure shown separately. Panels (c, d) use the first-condition held-feedback Curvature decision with the largest absolute nominated pitch target. Sampled derivatives describe geometric tightness and do not replace the analytic certificate

5 Discussion

The governor combines the path fraction to execute with the stroke reserve needed to stop. Directional slope and curvature bounds determine the motion envelopes; reconstruction error and response uncertainty determine the reserves. The finite-stop supervisor connects these calculations across successive updates, providing a model-specific construction of admissible motion and stopping within the reference-governor framework [6, 10].

The geometric ablation explains why the baseline matters. Attachment components parallel to the rotation axis contribute to the whole-vector norm but not to rotational motion. Removing them accounts for most of the recovered allowance, especially in Roll. Curvature adds a further increment by using the radial directional slope instead of the attachment-motion magnitude. The shared-input comparison with AxisConservative identifies that increment. The targeted closed-loop comparison measures its effect in the coupled platform, where native scheduling and feedback allow trajectories to diverge. It therefore estimates total execution behavior under matched configurations.

Certificate and search resolution should be chosen together. Curvature evaluates an analytic enclosure over the candidate prefix. DirectSearch subdivides that prefix to tighten local bounds, adding evaluations at each bisection step. Finer subdivision admits higher mean certified progress in the tested bank but requires more computation. The progress and decision-time distributions in Section 4.1 quantify this choice for the tested implementation and host, without implying a fixed latency ratio elsewhere.

The pitch response distinguishes unused actuator headroom from admissible motion under the response envelope. Speed and stroke margins do not release a decision whose desired-acceleration budget is active. Every eligible decision in the representative pitch phase selects pitch, and removing individual length reserves leaves the interior decision unchanged. Together with the sampled derivatives, these diagnostics identify the robust acceleration envelope and repeated zero-speed segments as the local explanation for attenuation. The reserves instead protect boundary cases when initial error or stopping displacement acts toward a limit. Full-amplitude reproduction requires a different response-budget and time-allocation design from the partial terrain presentation evaluated here.

Information-age modulation scales the certified path fraction, whereas extrapolation changes the state supplied to center selection. Better consumed-state estimates therefore need not produce the same ordering in terrain-tracking error. The feedback comparisons distinguish these functions within the same actuator supervisor; they do not establish a delay-compensation guarantee.

The admissibility result applies to the declared virtual actuator limits and response-time set; pose compatibility is checked separately through the reconstruction residual. The five model–seed conditions on Standard terrain define the closed-loop evidence, with feedback modes paired within each condition. AxisCurvature and the search-resolution scan retain their offline scope. Physical implementation requires identified actuator parameters and validation of the coupled platform response. The resting-set invariant preserves an admissible action, including a hold; positive progress and tracking convergence require additional viability and controller conditions.

6 Conclusion

The proposed governor uses directional actuator-length bounds to select a pose-path fraction while reserving stroke for initial error, uncertain response, and finite stopping. Bounded braking and a resting fallback preserve recursive actuator-constraint admissibility under the stated virtual model. In the shared-input benchmark, Curvature increases mean certified progress by 7.69% over AxisConservative at a 0.1-mm initial offset; AxisCurvature raises this offline gain to 9.73%. In the targeted closed-loop comparison, roll tracking

RMSE is 2.62%–3.70% lower than AxisConservative across ten configuration-matched paired comparisons. The decision-budget analysis explains why acceleration constraints can attenuate motion despite unused speed and stroke margins. The path benchmark and inherited-model specification support reproduction of the geometric comparisons and interpretation of the closed-loop behavior.

Appendix A Inherited closed-loop components

This appendix specifies the inherited environment [8] used in Figures 3–6. The added governor acts after center selection and before virtual-actuator integration.

A.1 Adaptive rotation-center map

Let $\bar{\phi}$ and $\bar{\psi}$ be the signed realized platform angles read by the center controller, $\mathbf{c}_h$ the consumed COM-offset vector, S the consumed stability index, and $\boldsymbol{\omega}_{\mathrm{deg}}$ the consumed angular-rate vector in degrees/s. The model-scaled center factors are

$$\begin{aligned} \mathbf{T} &= 1.5(0.15\bar{\psi}/15^{\circ}, 0, -0.18\bar{\phi}/15^{\circ}), \\ \mathbf{U} &= 0.75(1-S)\mathbf{c}_h, \\ \mathbf{P} &= 1.5\left[0.005\boldsymbol{\omega}_{\mathrm{deg}} + 0.15\mathbf{c}_h + 0.3(0.7-S)_+\hat{\mathbf{c}}_h\right]. \end{aligned} \tag{11}$$

The inherited readout uses $\bar{\phi} = -\mathrm{wrap}_{180}(\mathrm{Euler}_x)$ and $\bar{\psi} = -\mathrm{wrap}_{180}(\mathrm{Euler}_z)$; plotted realized roll uses the opposite sign to $\bar{\psi}$. The COM terms in $\mathbf{P}$ apply when $\| \mathbf{c}_h \| > 0.001$ m; the angular term applies when $\| \boldsymbol{\omega}_{\mathrm{deg}} \| > 0.001$ degrees/s. The hat denotes a unit vector, with zero used below the COM threshold. Nine fuzzy rules map $\mathrm{clip}(|\bar{\phi}|/15^{\circ}, 0, 1)$, S, and angular magnitude normalized by 1 rad/s to weights (w_T, w_U, w_P). Online Resource 2, Section S15, gives the rule outputs, memberships, and smoothing update. With $\mathbf{G} = \mathbf{T}[2/(1+\exp(-3\| \mathbf{T} \|)) - 1]$, the inherited offset is

$$\begin{aligned} \mathbf{c}_{\mathrm{raw}} = \; & w_T\mathbf{T} + w_U\mathbf{U} + w_P\mathbf{P} + 0.15\mathbf{G} \\ & + 0.08\exp(-1.5\| \mathbf{U} \|^2)\mathbf{G} \\ & + 0.11\tanh(2\chi)(\mathbf{U}+\mathbf{P})/2, \end{aligned} \tag{12}$$

where χ is the cosine between $\mathbf{U}$ and $\mathbf{P}$, set to zero if either norm is at most 10^{-6}. The center filter clips horizontal offset to 0.4 m and applies one 0.8 scaling when its 35° surrogate angle test is exceeded. Native LateUpdate smooths the horizontal target with SmoothDamp time constant $1/4.5$ s, while center height follows the platform. These inherited filters precede the actuator-length checks in Sections 2.2–2.5. An arriving-state gate holds the previous center until a packet is available.

A.2 Digital-patient profiles

Each profile combines physiological input parameters with a posture-strategy asset. The common body model uses height 1.75 m, mass 70 kg, and static COM height 0.9975 m. Table 7 lists the parameters that distinguish the named configurations. These names and scores are software inputs.

Table 7. Inherited digital-profile settings. The physiological and posture-strategy response parameters enter separate model components; neither is the packet-delivery delay.

Parameter	Healthy	Mild	Moderate
Muscle-strength score (MMT scale)	5	4	3
Balance score (BBS scale)	52	42	35
Physiological response parameter (ms)	400	550	700
Posture-strategy response parameter (ms)	100	250	350
Ankle-to-mixed / mixed-to-hip angle (deg)	6 / 10	5 / 9	3 / 6
COM-variability multiplier	1.0	1.4	1.8
Derived angular-response multiplier	1.0	0.7	0.5
Derived COM-drift amplitude (mm)	0	16	32

The scene's digital-patient fallback type is HealthyAdult for all three configurations; the selected posture asset supplies the runtime overrides. Its angular multiplier is $1 - 0.6\mathrm{clip}[(d_r - 100)/300,0,1]$, with d_r in ms. The profile-specific COM drift is $(A\sin(0.3t),0,A\cos(0.39t))$, where $A = \max[0,(1.2 - k_{\mathrm{ICP}})0.04]$ m and k_{ICP} is 1.2, 0.8, or 0.4. Native patient updates are eligible at 25 Hz; stability and angular-rate smoothing use coefficient 0.8. BodyTiltController supplies the orientation quaternion and COM state to this layer. Section S15 specifies the remaining stability and noise mappings.

A.3 Standard terrain

Terrain time starts after the 15-s adaptation phase. The 90-s Standard asset stores 26 pitch and 15 roll keys. Pitch keys sample two 30-s, 8° sinusoidal cycles over 0–60 s, with zero endpoint tangents and a zero 60–90-s interval. Roll keys sample a 15-s, 5° cycle over 60–75 s at 1.25-s spacing, followed by a cubic-Hermite return to zero at 90 s. In particular, the stored 75-s roll tangent remains approximately 2.0943952 degrees/s, producing a positive tail rather than a second sinusoidal cycle. Between adjacent keys (t_j, y_j, m_j^+) and $(t_{j+1}, y_{j+1}, m_{j+1}^-)$, the unweighted interpolation is

$$y(t) = h_{00}(u)y_j + h_{10}(u)\Delta t\, m_j^+ + h_{01}(u)y_{j+1} + h_{11}(u)\Delta t\, m_{j+1}^-, \qquad (13)$$

where $u = (t - t_j)/\Delta t$, $h_{00} = 2u^3 - 3u^2 + 1$, $h_{10} = u^3 - 2u^2 + u$, $h_{01} = -2u^3 + 3u^2$, and $h_{11} = u^3 - u^2$. Section S15 records the key-generation values and endpoint rules. The provider samples this interpolant in native Update, so realized acquisition times account for the small request differences between configuration-matched executions.

References

1. Barton GJ, Vanrenterghem J, Lees A, Lake M (2006) A method for manipulating a movable platform's axes of rotation: a novel use of the CAREN system. Gait & Posture 24(4):510–514. https://doi.org/10.1016/j.gaitpost.2005.12.011

2. Lees A, Vanrenterghem J, Barton G, Lake M (2007) Kinematic response characteristics of the CAREN moving platform system for use in posture and balance research. Medical Engineering & Physics 29(5):629–635. https://doi.org/10.1016/j.medengphy.2006.06.004

3. Garrett NJI, Best MC (2013) Model predictive driving simulator motion cueing algorithm with actuator-based constraints. Vehicle System Dynamics 51(8):1151–1172. https://doi.org/10.1080/00423114.2013.783219

4. Ellensohn F, Schwienbacher M, Venrooij J, Rixen D (2018) Motion cueing algorithm for a 9 DoF driving simulator: MPC with linearized actuator constraints. SAE Technical Paper 2018-01-0570, SAE International. https://doi.org/10.4271/2018-01-0570

5. Khusro YR, Zheng Y, Grottoli M, Shyrokau B (2020) MPC-based motion-cueing algorithm for a 6-DOF driving simulator with actuator constraints. Vehicles 2(4):625–647. https://doi.org/10.3390/vehicles2040036

6. Garone E, Di Cairano S, Kolmanovsky I (2017) Reference and command governors for systems with constraints: a survey on theory and applications. Automatica 75:306–328. https://doi.org/10.1016/j.automatica.2016.08.013

7. Jiang X, Wang C, Zeng H, Zhang B, Zhang L (2026) Online trajectory planning for Stewart parallel mechanisms based on improved trajectory scaling and condition-number-aware velocity bound regulation. Machines 14(7):794. https://doi.org/10.3390/machines14070794

8. Wang X, Niu Z, Zhang P, Pan J (2026) Adaptive center of rotation control for VR gait rehabilitation platform (in Chinese). Journal of System Simulation. Published online 17 June 2026.

9. Kassis T, Agarwal V, He Y, Patel D, Brueckner AM (2026) Scientific Agent Skills: A Library of Procedural Knowledge for Research Agents. arXiv:2609.00065. https://doi.org/10.48550/arXiv.2609.00065

10. Gilbert E, Kolmanovsky I (2002) Nonlinear tracking control in the presence of state and control constraints: a generalized reference governor. Automatica 38(12):2063–2073. https://doi.org/10.1016/S0005-1098(02)00135-8

11. Di Cairano S, Kalabić UV, Kolmanovsky IV (2015) Reference governor for network control systems subject to variable time-delay. Automatica 62:77–86. https://doi.org/10.1016/j.automatica.2015.09.006

12. Pham H, Pham QC (2018) A new approach to time-optimal path parameterization based on reachability analysis. IEEE Transactions on Robotics 34(3):645–659. https://doi.org/10.1109/TRO.2018.2819195