

When is a rigid pelvis–trunk approximation sufficient? An output-specific model-reduction protocol applied to lumbar mobility in multibody biomechanics

Toufic Antoun El Halabi^{1*} and Nathalie Chahine²

¹ Physics Department, Faculty of Science – Branch 2, Lebanese University, Lebanon

² Faculty of Public Health– Branch 2, Lebanese University, Lebanon

* Corresponding author: tofic.anton@hotmail.com

ORCID: T.A. El Halabi <https://orcid.org/0009-0000-5409-0417>;

N. Chahine <https://orcid.org/0000-0002-6200-7617>

Abstract

Simplified human multibody models often lump the pelvis and trunk into one rigid body, but model reduction is defensible only for outputs and operating ranges that it preserves. We compare exact lumbar locking, axial-yaw-only mobility and three-coordinate mobility in a support-driven single-leg task using the public ModelicaHumanBodyPArts v0.10.0 release. The campaign combines mass-equivalent architecture ablation, an analytical benchmark, mirrored-side checks, numerical refinement, a 15-condition operating sweep, lumbar-impedance sensitivity and a rigid upper-body-inertia sensitivity. The primary criterion is a 10% locked-normalized maximum-error ratio, supplemented by RMS and actuation-window audits. With nominal transverse stabilization (1000 N·m/rad), yaw-only mobility met the criterion for pelvis and thorax yaw in all 15 primary conditions, whereas horizontal ground force failed in all 15. At 100 N·m/rad transverse stiffness, every audited output failed, including pelvis yaw (30.4%) and thorax yaw (17.2%). Restoring rigid head and upper-limb inertia met the same 10% criterion for thorax yaw in 15/15 conditions and pelvis yaw in 13/15, while several support outputs became substantially less transferable. The resulting validity boundary shows that reduction adequacy depends jointly on output, transverse stabilization, inertial scope and operating condition. These results verify a computational reduction protocol, not a human or clinical model.

Keywords: multibody dynamics; model reduction; lumbar spine; model verification; support mechanics; inertia matching; joint moments.

1. Introduction

A simplified human model must preserve the mechanisms relevant to its intended question. A rigid pelvis–trunk representation can be appropriate when relative axial motion contributes little to the outputs of interest. It cannot, however, represent a relative lumbar response by definition. Deciding whether that missing motion matters requires a controlled comparison, not the assumption that a more articulated model is automatically better.

Human models with lumbar articulation already exist. General-purpose musculoskeletal platforms such as OpenSim (Delp et al. 2007; Seth et al. 2018) have been used to build fully articulated thoracolumbar models (Bruno et al. 2015) and detailed lumbar models validated for lifting tasks (Beaucage-Gauvreau et al. 2019). Raabe and Chaudhari (2016) developed and validated a full-body lumbar-spine model for jogging, and Shin et al. (2013) reported coupled bending during dynamic axial rotation. Parametric multibody human models also make explicit choices about segment partition, joint architecture and inertial properties (Bova et al. 2020). More detailed lumbar models include the text-based rigid-body model of de Zee et al. (2007) and the three-rotation lumbar model of Christophy et al. (2012). A review by Lerchl et al. (2023) shows that thoracolumbar multibody modeling already spans many levels of detail and repeatedly faces the trade-off between necessary accuracy and manageable complexity. Recent rigid-flexible work further illustrates the continuing search for computationally efficient spinal representations (Dukkipati and Driscoll 2026).

More generally, computational-biomechanics guidance treats model credibility as purpose- and output-specific rather than as a single global property. Verification and validation frameworks emphasize intended use, the outputs being predicted, sensitivity to modeling assumptions and comparison with independent evidence (Henninger et al. 2010; Lund et al. 2012; Hicks et al. 2015). Classical model-order reduction in multibody dynamics is mathematically different from the architecture ablation used here, but it likewise evaluates reduced models by how well specified input-output behavior is preserved over a declared domain (Lehner and Eberhard 2006; Vierendeis et al. 2026). Joint-model studies similarly show that increasing mechanical detail does not uniformly improve every prediction and that appropriate complexity depends on the task and intended output (Valente et al. 2015). The present study adopts that logic explicitly: the question is not whether one lumbar architecture is globally 'better,' but where a simpler architecture remains adequate for named observables under controlled assumptions.

A more specific question remains useful for the present application: when total mass and reference-pose inertia are controlled, how much does permitting lumbar motion change support-to-trunk transmission, heel support and the opposite-leg response? Replacing an old lower-trunk proxy with a different pelvis and thorax at the same time as adding a joint confounds mass changes with mobility changes. A comparison that avoids that confounding can identify the range in which a simpler architecture is sufficient.

The motivating task, described physically in Section 2.4, is single-leg support in which efforts are declared at the standing ankle and hip while pelvis, trunk and free-leg responses are left as outputs. No preferred outcome is assumed: greater free-side motion, lower sustained moments, higher local peaks, no measurable change and loss of support are all possible outcomes.

The contribution is an auditable output-specific reduction procedure: mass-equivalent architecture ablation, an independent known-answer benchmark, mirrored-side verification, numerical refinement, parameter and inertial sensitivity, and explicit error criteria. The lumbar joint is the worked example. The procedure is intended as a transparent decision framework for constrained multibody subsystems, not as evidence that one lumbar architecture is universally superior.

2. Model and methods

2.1. Common mechanical body

All comparisons use the same bilateral lower limbs. Each side contains rigid thigh, shank and foot segments; three hip rotations; one knee rotation; and two ankle rotations. The axial body contains a rigid pelvic ring carrying both hip connections, an effective lumbar joint, and a rigid upper-trunk segment. The pelvic ring is not divided into freely rotating left and right halves. Sacroiliac compliance and individual vertebral motion are outside the baseline model.

Segment parameters are supplied through explicit mass, center-of-mass and inertia records. Selected lower-limb coefficients follow published sources, while axial properties require a declared partition (Bova et al. 2020; de Leva 1996). The baseline fixture represents 62.79 kg and omits explicit head, upper-limb and articulated-neck bodies. A separate sensitivity campaign rigidly restores de Leva-based head and upper-limb inertia to the thoracic/cervical frame; it is treated as an inertial-scope test rather than a physiological arm/shoulder model. Within every A/B/C comparison, the same represented segments and initial pose are retained and only lumbar constraints change.

The same geometry, feet, contact law, initial task direction and standing-side actuator channels were used in every architecture comparison. Identical solver settings were used for the nominal cases, and every reported case was required to pass the same runtime and balance checks. The controlled comparison uses synthetic and source-derived parameters whose status is declared; it is not presented as physiological calibration.

2.2. Three axial architectures

Table 1. Primary architecture factor. A very large spring is not silently substituted for an exact lock.

Architecture	Lumbar motion	How the comparison is controlled
A: locked	No relative lumbar rotation	Retain the same pelvic and trunk masses; apply explicit constraints
B: yaw-mobile	Axial rotation only	Release only the axial coordinate; keep the other two constrained
C: three-coordinate	Flexion, lateral bending and axial rotation	Release all three coordinates with declared effective impedance coefficients

The same bilateral hip joints were retained in all three architectures. Effective lumbar stiffness and damping coefficients for each available coordinate were held constant in the nominal comparison, and time-varying impedance was excluded so that architecture remained the primary factor. Nominal values and their status are listed in Table 2. A separate coefficient-sensitivity campaign varied transverse and axial stiffness and damping. The axial stiffness sweep of 10–40 N·m/rad (0.17–0.70 N·m/deg) overlaps the early-to-middle weighted in-vivo axial-rotation stiffness range summarized by Watt et al. (2022), whereas the transverse sweep

was deliberately broader to test the reduction boundary rather than to represent passive tissue alone.

Table 2. Nominal effective lumbar stiffness and damping, sensitivity ranges and status relative to published in-vivo data.

Coordinate	Nominal k [N·m/rad]	Nominal c [N·m·s/rad]	Sensitivity range	Physiological reference	Status
Axial rotation (+z)	20	1	k: 10–40 N·m/rad c: 0.5–2 N·m·s/rad	Passive in-vivo axial stiffness (Watt et al. 2022)	Physiologically informed
Flexion/extension (+y)	1000	50	k: 100–3000 N·m/rad c: 10–100 N·m·s/rad	In-vivo passive flexion ≈ 0.3 –3 N·m/deg from early to late range (Watt et al. 2022)	Lower end within late-range values; nominal and upper end: surrogate for active stabilisation
Lateral bending (–x)	1000	50	k: 100–3000 N·m/rad c: 10–100 N·m·s/rad	In-vivo passive lateral bending ≈ 0.3 –5 N·m/deg from early to late range (Watt et al. 2022)	Lower end within late-range values; nominal and upper end: surrogate for active stabilisation

The transverse stiffness range deserves comment. Its lower end, 100 N·m/rad (≈ 1.7 N·m/deg), is comparable to stiffer late-range passive values reported in vivo for flexion and lateral bending, but it should not be interpreted as neutral-zone stiffness. The nominal 1000 N·m/rad (≈ 17 N·m/deg) and upper 3000 N·m/rad (≈ 52 N·m/deg) values are supra-passive. Because the model contains no explicit trunk musculature, they are used as lumped effective-stabilization surrogates, consistent with the distinction between passive spinal tissues and active stabilizing contributions (Panjabi 1992a). The sensitivity sweep therefore tests whether the reduction conclusion survives progressive relaxation of that effective stabilization rather than claiming that the nominal coefficient is a passive tissue property.

For a permitted rotational coordinate,

$$Q_p = -k(q - q_0) - c \dot{q}, \quad (1)$$

where Q_p is passive generalized torque, k is rotational stiffness, c is rotational damping, q is the relative coordinate, and q_0 is the declared neutral coordinate. The library basis is +x anterior, +y left and +z superior. The effective lumbar coordinates are applied in serial order: flexion/extension about +y, lateral bending about –x, then axial rotation about +z.

A constrained coordinate transmits the reaction needed to maintain the constraint. In an ideal stationary relative lock it does no relative joint work. This does not mean that its reaction moment is zero, or that a real person could maintain the same posture without muscular effort. Figure 1 depicts the three architectures while retaining the same pelvic and trunk partitions and bilateral hips.

Controlled lumbar architectures

Same pelvis, thorax, hips, lower limbs, contact and reference mass properties in all three models.




Figure 1. Axial architectures compared in the study. A locks the lumbar connection, B permits axial rotation only, and C permits flexion, lateral bending and axial rotation. The pelvic and trunk rigid bodies are identical in all three models.

2.3. Reference-pose mass equivalence

Before dynamic comparison, the three architectures were constructed from the same represented segments and were checked to reproduce the same total mass and reference-pose center of mass. The reference-pose inertia was evaluated about a common origin and basis:

$$M = \sum_i m_i, \quad r_C = (\sum_i m_i r_C,i) / M. \quad (2)$$

The inertia tensor about a common origin O is

$$I_O = \sum_i [R_i I_C,i R_i^T + m_i ((d_i^T d_i) 1 - d_i d_i^T)], \quad (3)$$

where $d_i = r_C,i - r_O$ is the vector from the common origin O to segment i's center of mass, and R_i rotates a segment's inertia basis into the common basis. Equation (3) combines rotation of the center-of-mass tensor with the parallel-axis theorem. Matching only total mass is insufficient: a different center of mass or rotational inertia could explain a different response.

No separate rigid replacement body is introduced in the primary A/B/C experiment. The same pelvis, thorax and lower-limb segments are instantiated in all architectures; the locked case constrains all three lumbar rotations exactly, the yaw-only case releases only axial rotation, and the three-coordinate case releases all three rotations. Consequently total represented mass, reference-pose center of mass and reference-pose inertia are identical across A/B/C by construction. This avoids the confound in which 'adding a spine' simultaneously changes body partition or inertia. Figure 2 summarizes the mass-property quantities controlled before dynamic differences are interpreted as architecture effects.

Mass-controlled comparison | equivalence at one reference pose




Figure 2. Reference-pose mass-equivalence requirements. The same segment masses, centers of mass and center-of-mass inertia tensors are retained. After transformation to a common reporting frame and origin, total mass, whole-body center of mass and full inertia are controlled before dynamic comparison.

2.4. Dynamics, task and declared effort

The assembled equations can be written as

$$M(q) \ddot{q} + h(q, \dot{q}) = S^T \tau_a + Q_p + J_c^T \lambda + J_o^T w_o. \quad (4)$$

The assembled equations use the standard multibody form in which the mass matrix, gravity/velocity terms, declared actuator torques, passive generalized efforts and contact wrenches are solved together. The implementation uses the component equations of the Modelica MultiBody library (Otter et al. 2003) rather than an independently invented dynamics engine.

The primary trial drives the standing ankle pitch and standing hip yaw channels. It does not command the free leg or prescribe pelvis and lumbar trajectories. In the coupled system, an ankle actuator need not produce pure ground shear, and a hip actuator need not produce pure ground yaw. Results therefore retain actuator-channel names alongside the actual measured contact wrench.

Physically, the simulated task corresponds to single-leg stance on the right foot (or the reflected left foot), during which a brief plantarflexion effort at the standing ankle and a small rotational effort at the standing hip are applied through a smooth pulse beginning at 0.25 s, with a 0.10 s rise, 0.60 s hold and 0.20 s fall (0.90 s total), while the contralateral limb is unloaded. The resulting motion is a small, transient rotation of the pelvis and trunk about the vertical axis, accompanied by a forward excursion of the free limb. It resembles in kind, though not in amplitude or physiological actuation, a supported turning preparation.

2.5. Independent rotational benchmark

An independent known-answer control isolated relative axial mechanics before the gravity-loaded body comparison. It used two rotational inertias, $J_p = 1.2 \text{ kg}\cdot\text{m}^2$ and $J_t = 0.8 \text{ kg}\cdot\text{m}^2$,

coupled by a torsional spring $k = 15 \text{ N}\cdot\text{m}/\text{rad}$ and released from a relative angle $\delta_0 = 0.10 \text{ rad}$ with zero initial angular velocities and zero external torque.

With relative angle $q = \theta_p - \theta_t$ and positive external torque u acting on the pelvis, Eqs. (5)–(6) give the coupled two-inertia relation:

$$J_p \ddot{\theta}_p = u - kq - c \dot{q}, \quad J_t \ddot{\theta}_t = kq + c \dot{q}, \quad (5)$$

$$\ddot{q} + c(1/J_p + 1/J_t) \dot{q} + k(1/J_p + 1/J_t) q = u/J_p. \quad (6)$$

The executed benchmark used $c = 0$ and $u = 0$, reducing the relative coordinate to the closed-form undamped mode $\delta(t) = \delta_0 \cos(\omega_n t)$, with $\delta \equiv q$ and

$$\omega_n = \sqrt{k(1/J_p + 1/J_t)}. \quad (7)$$

For the declared parameters, $\omega_n = 5.59017 \text{ rad/s}$. Conservation of zero total angular momentum gives $\theta_p = J_t/(J_p + J_t) \cdot \delta$ and $\theta_t = -J_p/(J_p + J_t) \cdot \delta$, i.e. $\theta_p(0) = 0.04 \text{ rad}$ and $\theta_t(0) = -0.06 \text{ rad}$. The native OpenModelica run uses 2000 output intervals over 1 s with tolerance 10^{-10} and is compared point-by-point with the closed-form solution, using an acceptance bound of $2 \times 10^{-6} \text{ rad}$ on δ , θ_p and θ_t . Results are reported in Section 3.1.

2.6. Load-bearing foot contact and initialization

The functional trial used load-bearing contact components on both feet. The standing foot was not held by a parallel fixed support: its reaction arose from the contact model and body dynamics, while the free foot remained unloaded until its geometry created contact. The legacy fixed-foot experiment was retained only as historical background and was not used as evidence for the present architecture comparison.

The contact formulation was required to carry the intended planted-foot effort. Because the velocity-regularized friction law permits finite creep, stance-foot slip was measured rather than assumed to be zero. Normal loading, shear, support location and yaw capacity were treated as coupled outputs rather than independent unlimited controls (Caron et al. 2015).

Each architecture began from a gravity-loaded initialized state using the same root pose, joint coordinates, contact parameters and input schedule. No balance controller and no temporary support was active during the analysis interval; the root was free and the stance reaction was generated by contact. Initial values used for the nominal comparison are reported in Table 3.

The common primary initialization preserved the same represented mass, segment configuration, heel share, broad-support state and stance-axis angle. Exact locking can transmit a different initial lumbar reaction from a mobile joint, so initial cut reactions were reported rather than forced to match. The yaw-only axial spring itself starts at zero generalized torque ($q = q_0$ and $\dot{q} = 0$); small world-frame cut-moment differences arise from the resolved three-dimensional reaction state and are not treated as an imposed axial preload.

Table 3. Nominal right-stance initialization. Left stance used the reflected configuration; scalar support measures are unchanged and axial reaction signs reflect.

Architecture	Represented mass [kg]	Pelvis/lumbar/thorax yaw [rad]	Heel share	Stance axis [rad]	Initial lumbar Mz [N·m]
Locked	62.79	0 / 0 / 0	0.5963	-0.004326	-0.01936
Yaw only	62.79	0 / 0 / 0	0.5963	-0.004326	-0.001041
3-D lumbar	62.79	0 / 0 / 0	0.5963	-0.004326	≈ 0

2.7. Right and left stance conventions

World axes are +x forward, +y left and +z upward. Clockwise and anticlockwise are defined from above, looking down at the floor. In the right-standing drive phase, the intended foot-on-floor action is backward shear and clockwise torsional effort. The floor-on-foot reaction has the opposite signs. In the left-standing reflection, backward shear retains its task direction and the world yaw sign reverses. These effort directions are not imposed foot rotations. For reflection matrix $S = \text{diag}(1, -1, 1)$,

$$F' = S F, \quad M' = \det(S) S M. \quad (8)$$

Moment is an axial vector, so its reflection differs from that of force. Segment geometry, side roles, preload, actuation and regional foot labels are reflected together. This is a verification control for symmetric parameters, not an assumption that individual human subjects are symmetric.

2.8. Experiment matrix

The nominal series contained six cases: three axial architectures $\times$ two stance sides, all using the same combined standing-side drive, initial configuration, contact law and represented mass. The right-stance cases formed the primary mechanical comparison and the reflected left-stance cases checked sign symmetry. The robustness campaign contained 15 operating conditions $\times$ three architectures (45 simulations), the lumbar-impedance campaign contained 10 coefficient settings $\times$ two mobile architectures (20 simulations), and numerical refinement covered both sides and all three architectures at three integration settings (18 runs). An additional upper-body inertia campaign used the same 15 operating conditions across all three architectures (45 simulations), following a known-answer inertia check and three preliminary nominal audit runs (Table 4).

Table 4. Executed campaign design and controlled factors.

Study block	Executed design
Nominal architecture comparison	3 architectures $\times$ 2 stance sides = 6 simulations
Operating robustness	15 conditions (Table 5) $\times$ 3 architectures = 45 simulations

Study block	Executed design
Lumbar-impedance sensitivity	10 coefficient settings $\times$ yaw-only / 3-D = 20 simulations
Numerical refinement	2 stance sides $\times$ 3 architectures $\times$ 3 numerical settings = 18 simulations
Upper-body inertia sensitivity	Head and upper-limb mass properties (de Leva 1996) rigidly attached to the thorax, with no active upper-limb torque; known-answer inertia check + 15 conditions $\times$ 3 architectures = 45 simulations (plus 3 preliminary nominal audit runs)
Common mechanics	Same segmented mass properties, contact formulation and declared initial configuration
Primary outputs	Pelvis/lumbar/thorax yaw, free-side x, heel share, stance axis, ground wrench and lumbar cut moment

The nominal input applied a 15 N·m standing-ankle channel and a -1 N·m standing-hip-yaw channel with identical pulse timing in all three architectures. These inputs are synthetic mechanical loads rather than physiological constants. The operating sweep additionally covered ankle amplitudes of 10, 15 and 20 N·m, hip-yaw amplitudes of -0.5 , -1 and -1.5 N·m, two friction values, two contact stiffnesses and two initial free-ankle postures. The 15 conditions retained from these levels are defined in Table 5.

For the supplementary upper-body-inertia campaign, the baseline lower-body rest angles and contact geometry were deliberately left unchanged so that the perturbation remained an inertial-scope test. This means the original contact indentation initially supports the 62.79-kg baseline rather than the 75.405-kg supplemented mass, producing a short unforced settling transient before the actuator pulse begins at 0.25 s. A post-hoc audit therefore evaluates the supplementary campaign both over the full 0–1.5 s interval and over the driven window $t \geq 0.25$ s. At $t = 0.25$ s, vertical ground reaction was approximately 740.5–740.7 N versus 739.72 N body weight, vertical whole-body COM acceleration was approximately 0.013 m/s², and root vertical speed was on the order of 10^{-5} m/s, indicating that the initial vertical settling was effectively complete before the imposed pulse. No lower-body coefficient was retuned for this audit.

Table 5. Definition of the 15 operating conditions in the primary robustness sweep. HGF = horizontal ground force; ground Mz = ground yaw moment; lumbar Mz = lumbar axial cut moment. The last column lists outputs for which the yaw-only model did not meet the 10% criterion.

#	Ankle [N·m]	Hip yaw [N·m]	Friction μ	Contact k_n [N/m]	Free ankle q_0 [rad]	Outputs failing $\eta \leq 10\%$
1	15 (nominal)	-1.0	0.60	200000	-0.03	HGF
2	10	-0.5	0.60	200000	-0.03	HGF
3	10	-1.0	0.60	200000	-0.03	HGF
4	10	-1.5	0.60	200000	-0.03	HGF
5	15	-0.5	0.60	200000	-0.03	Free-side x; HGF
6	15	-1.5	0.60	200000	-0.03	HGF
7	20	-0.5	0.60	200000	-0.03	Free-side x; heel share; HGF
8	20	-1.0	0.60	200000	-0.03	HGF
9	20	-1.5	0.60	200000	-0.03	HGF
10	15	-1.0	0.45	200000	-0.03	HGF
11	15	-1.0	0.75	200000	-0.03	HGF
12	15	-1.0	0.60	100000	-0.03	Heel share; HGF; ground Mz
13	15	-1.0	0.60	400000	-0.03	Heel share; stance axis; HGF; ground Mz; lumbar Mz
14	15	-1.0	0.60	200000	-0.02	HGF
15	15	-1.0	0.60	200000	0.00	HGF

Identical input histories were used so that different trajectories could be attributed to architecture and operating parameters. The study therefore compares model responses under matched inputs; it does not claim equal-task energetic optimality or human movement performance.

2.9. Outcome measures

Primary outputs were pelvis and thorax yaw, relative lumbar yaw, free-side horizontal response, stance-axis response, ground yaw moment, and the lumbosacral cut moment. Resultant lumbosacral moment exposure was defined by Eq. (9) as the time integral of the magnitude of the three-component cut moment:

$$\mathcal{E}_{M,j} = \int_{\{t_0\}^{\wedge}\{t_f\}} \|M_j(t)\| dt, \quad (9)$$

with units N·m·s. This exposure is not energy, muscle activation or a damage measure. The reporting basis and cut origin were fixed across architectures. The archive also retains hip, knee, ankle and regional plantar channels for audit, but those channels were not used to classify the model-reduction adequacy reported here.

Secondary outputs used in the reduction analysis were heel share, broad-support status and selected ground-reaction quantities. Free-side position was evaluated relative to the moving hip/pelvis so that a change in body orientation was not mistaken for leg extension. Heel share and broad support were interpreted separately. A high heel fraction caused by forefoot unloading was not treated as evidence of a larger contact area, and regional contact quantities were interpreted only when the load-bearing denominator remained valid. Figure 3 separates absolute segment orientation, relative lumbar motion, cut reactions and the secondary support/free-side measurements.




Figure 3. Output map used for reduction testing. Absolute pelvis and thorax orientation are separated from relative lumbar motion. The lumbosacral cut moment, ground reaction, support measures and free-side response are reported in a fixed basis.

2.10. Verification, reduction metrics and sensitivity

Before architecture effects were interpreted, the nominal cases were checked for the same segmented body, represented mass, initial configuration and contact formulation; reflected left/right cases were used as a sign-symmetry check. Each nominal side/architecture case was then evaluated at three numerical settings (base, fine and finer), producing 18 refinement runs.

Sensitivity was evaluated in three complementary campaigns. The 15-condition operating sweep varied drive level, friction, contact normal stiffness and free-ankle posture across all three architectures. The lumbar-impedance sweep varied transverse stiffness, transverse damping, axial stiffness and axial damping while comparing yaw-only with three-coordinate mobility. The upper-body inertia campaign restored rigid head and upper-limb mass properties without active upper-limb torque. To isolate inertial scope, that supplementary campaign retained the baseline geometry, joint rest parameters, contact parameters and root/joint initial

coordinates rather than solving a new 75-kg static support state; re-equilibration would itself change support preload and contact state. Because the added mass therefore produces a short unforced settling transient before actuator onset, all upper-body conclusions were audited both over the full interval and again from $t \geq 0.25$ s. Two additional post-hoc audits were performed on the frozen trajectories without changing model source or raw simulation results: an RMS error-ratio audit, and the same actuation-window audit beginning at $t = 0.25$ s to distinguish pre-actuation contact settling from the driven response. The transverse-stiffness audit was also expanded from selected kinematic plots to the full output set at 100, 300, 1000 and 3000 N·m/rad for the nominal condition.

For the final analysis, architectures were labeled A = locked, B = yaw-only and C = three-coordinate. For each selected output y and operating condition, $E_A = \max_t |y_A(t) - y_C(t)|$ and $E_B = \max_t |y_B(t) - y_C(t)|$ were computed on the common time base. The reported locked-normalized error ratio was $\eta_y = E_B / \max(E_A, 10^{-15})$, where the small denominator guard only prevents division by numerical zero. Thus $\eta = 100\%$ means that yaw-only is no closer to the three-coordinate reference than the locked model, whereas $\eta \leq 10\%$ means that its maximum deviation is at most 10% of the locked-model deviation. The 10% value is a stated engineering reporting criterion, not a preregistered biological or clinical threshold. For the primary operating-range classification, an output was accepted only where η remained $\leq 10\%$ and numerical-refinement error was negligible relative to the architecture difference. The transverse-stiffness and upper-body campaigns are reported as sensitivity analyses using the same η threshold after their stated verification gates; their pass counts are not presented as independently refined primary classifications. Deterministic time samples were not treated as independent human observations.

Because η can be inflated when the locked model is itself close to the reference, absolute deviations E_B are reported alongside η . Threshold sensitivity repeats the classification at 5% and 20% on the same trajectories. A secondary RMS ratio, $\eta_{\text{RMS}} = \text{RMS}(y_B - y_C) / \max[\text{RMS}(y_A - y_C), 10^{-15}]$, tests whether conclusions depend on a single transient maximum; it is reported as a robustness measure, not as a replacement primary criterion. The actuation-window audit recomputes the same maximum and RMS ratios for $t \geq 0.25$ s. No independent absolute-error pass rule is imposed because output-specific practical tolerances have not been experimentally validated for this synthetic task.

Figure 4 states the output-specific decision rule. The same reduced architecture can therefore be adequate for rotational kinematics and inadequate for a contact-force output under the same operating condition.




Figure 4. Output-specific model-reduction rule. For each output, the maximum yaw-only deviation from the three-coordinate reference is divided by the corresponding locked-model deviation. A pass requires $\eta \leq 10\%$ after verification and numerical refinement. The threshold is an engineering reporting criterion, not a clinical or biological limit.

2.11. Software and source provenance

The reproducibility campaign loaded the public ModelicaHumanBodyParts v0.10.0 release (El Halabi 2026) and separate Article 2 experiment packages. The public release is the verified RC1 source promoted without code changes; its historical RC1 filename is retained for provenance, while the version-specific software archive is <https://doi.org/10.5281/zenodo.22941517>. The primary and supplementary experiment sources are preserved separately, and their SHA-256 identifiers are recorded in the archived source manifest rather than repeated here. Simulations used OpenModelica 1.27.1 (64-bit), Modelica Standard Library 4.1.0 and Windows x64. The supplementary known-answer test, force-balance and energy-ledger gates were verified independently of the manuscript analysis.

3. Results


3.1. Benchmark, mass equivalence and mechanical verification

The independent two-inertia benchmark passed its internal assertion and reproduced the closed-form undamped relative mode with maximum absolute errors of 7.57×10^{-10} rad for δ , 3.03×10^{-10} rad for θ_p and 4.54×10^{-10} rad for θ_t , several orders of magnitude below the 2×10^{-6} rad acceptance bound (Figure 5, Table 6).

(a)

Independent two-inertia torsional benchmark

Undamped relative mode with zero total angular momentum after release.



(b)




Figure 5. Independent two-inertia benchmark. (a) Two rotational inertias coupled by an undamped torsional spring and released from a declared relative angle with zero external torque. (b) Native OpenModelica relative angle versus the closed-form mode; the curves are visually coincident and the maximum absolute error is 7.57×10^{-10} rad.

Table 6. Selected native and closed-form values from the executed two-inertia benchmark.

t [s]	δ Modelica [rad]	δ analytical [rad]	error [rad]	Status
0.25	0.017238838535	0.017238838378	1.57×10^{-10}	PASS
0.50	-0.094056449685	-0.094056449027	6.58×10^{-10}	PASS
0.75	-0.049667317208	-0.049667316843	3.65×10^{-10}	PASS
1.00	0.076932312587	0.076932312072	5.14×10^{-10}	PASS

All six nominal architecture cases used the same v0.10.0 source, experiment package, segment records, represented mass (62.79 kg), initial pose, contact law and actuator histories; architecture changed only the permitted lumbar coordinates. Reference mass properties were identical across A/B/C by construction, left/right reflection reproduced the expected sign symmetry, and all 18 numerical-refinement runs completed successfully. Fine-to-finer differences were orders of magnitude smaller than the architecture effects used for classification. Re-execution against the public v0.10.0 source reproduced the archived compared signals exactly.

Table 7. Completed verification outcomes for the nominal architecture comparison.

Verification outcome	Locked	Yaw-mobile	Three-coordinate
Reference mass and common configuration	PASS — same segmented mass and configuration	PASS — same segmented mass and configuration	PASS — same segmented mass and configuration
Reference center of mass and inertia	PASS — r_C and I_O identical (by construction)	PASS — r_C and I_O identical (by construction)	PASS — r_C and I_O identical (by construction)
Mirror/sign agreement	PASS — reflected left/right yaw signs	PASS — reflected left/right yaw signs	PASS — reflected left/right yaw signs
Dynamic contact and refinement	PASS; both sides refined at base/fine/finer settings	PASS; both sides refined at base/fine/finer settings	PASS; both sides refined at base/fine/finer settings

3.2. Controlled architecture effects

All six nominal cases passed the runtime checks and retained broad support through the 1.5 s analysis interval. Right/left mirroring reproduced yaw magnitudes with sign reversal to numerical precision. Table 8 reports the principal nominal outputs.

Under this particular synthetic task, allowing axial lumbar motion changed the distribution of pelvis and thorax yaw and increased the free-side forward response relative to the locked case (Figures 6 and 7). The yaw-only and three-coordinate architectures were close to one another, so the present result does not show a large added benefit from the two extra lumbar coordinates for these chosen outputs.




Figure 6. Right-stance nominal architecture comparison. The locked model forces pelvis and thorax yaw to coincide; the mobile architectures permit relative axial motion.




Figure 7. Right-stance nominal free-side response at 1.5 s. Yaw-only and three-coordinate lumbar models are similar and both differ from the locked case.

Table 8. Nominal architecture results from the v0.10.0 reproducibility campaign.

Stance	Architecture	Pelvis yaw final [rad]	Lumbar yaw final [rad]	Thorax yaw final [rad]	Free-side x final [mm]	Min heel share
Right	Locked	-0.03055	0	-0.03055	1.83	0.498
Right	Yaw only	-0.00694	+0.00825	+0.00126	10.86	0.491
Right	3-D lumbar	-0.00662	+0.00930	+0.00266	10.13	0.492
Left	Locked	+0.03055	0	+0.03055	1.83	0.498

Stance	Architecture	Pelvis yaw final [rad]	Lumbar yaw final [rad]	Thorax yaw final [rad]	Free-side x final [mm]	Min heel share
Left	Yaw only	+0.00694	-0.00825	-0.00126	10.86	0.491
Left	3-D lumbar	+0.00662	-0.00930	-0.00266	10.13	0.492

For right stance, the yaw-only model differed from the three-coordinate model by 0.000324 rad in final pelvis yaw, 0.00105 rad in final lumbar yaw, 0.00140 rad in final thorax yaw and 0.739 mm in final free-side horizontal response. By contrast, the locked model forced pelvis and thorax yaw to remain identical and produced only 1.83 mm of final free-side response versus 10.13 mm for the three-coordinate model.

The nominal rotations are small: final pelvis and thorax yaw remained below 0.031 rad ($\approx 1.8^\circ$) and relative lumbar yaw below 0.01 rad ($\approx 0.5^\circ$). For comparison, lumbar axial rotation is commonly reported on the order of degrees per motion segment, with total available motion depending on level and loading (White and Panjabi 1990). This comparison is used only to place the simulated amplitudes in context; the linear spring representation is treated here as a local effective approximation and not as a validated constitutive model of lumbar tissues.

The primary lumbosacral moment-exposure result was small: $\mathcal{E}_M$ was 10.742 N·m·s for the locked model, 10.686 N·m·s for yaw-only and 10.697 N·m·s for the three-coordinate model on right stance. The corresponding peak resultant lumbosacral moments were 7.519, 7.602 and 7.625 N·m. Thus the architecture strongly redistributed axial organization while changing the resultant moment exposure by less than 0.6% in this nominal task.

3.3. Validity of the rigid approximation

Applying the 10% output-specific criterion to the 15-condition primary sweep produced a mixed adequacy map rather than a single yes/no verdict. Over the full 0–1.5 s interval, yaw-only met the criterion in 15/15 conditions for pelvis yaw, 15/15 for thorax yaw, 13/15 for free-side displacement, 12/15 for heel share, 14/15 for stance-axis angle, 0/15 for horizontal ground force, 13/15 for ground yaw moment and 14/15 for axial lumbar cut moment. Median/max η for pelvis yaw were 4.94/7.29% and for thorax yaw 2.79/4.34%. Table 9 reports the complete map and absolute deviations.

Table 9. Output-specific yaw-only adequacy over the 15-condition operating sweep. η is the maximum yaw-only deviation from the three-coordinate reference divided by the corresponding locked-model deviation; a pass requires $\eta \leq 10\%$. The last column reports the median and maximum, across conditions, of $E_B = \max_t |y_{\text{yaw}} - y_{3D}|$ in the output's native units.

Output	Median η [%]	Maximum η [%]	Conditions $\eta \leq 10\%$	Median / max absolute deviation
Pelvis yaw	4.94	7.29	15 / 15	0.00118 / 0.00179 rad
Thorax yaw	2.79	4.34	15 / 15	0.00140 / 0.00235 rad

Output	Median η [%]	Maximum η [%]	Conditions $\eta \leq 10\%$	Median / max absolute deviation
Free-side horizontal response	7.41	11.01	13 / 15	0.738 / 1.046 mm
Heel share	7.11	101.22	12 / 15	0.00350 / 0.15188
Stance-axis angle	7.41	11.43	14 / 15	0.000316 / 0.000440 rad
Horizontal ground force	18.20	131.86	0 / 15	0.715 / 16.030 N
Ground yaw moment	3.96	68.97	13 / 15	0.0468 / 0.7650 N·m
Lumbar axial cut moment	2.55	10.64	14 / 15	0.0307 / 0.2379 N·m

The secondary audits separate persistent architecture sensitivity from isolated transients. Using η_{RMS} on the same full trajectories, pass counts were 15/15 for pelvis yaw, 15/15 for thorax yaw, 15/15 for free-side displacement, 11/15 for heel share, 12/15 for stance-axis angle, 0/15 for horizontal ground force, 14/15 for ground yaw moment and 15/15 for axial lumbar cut moment. Restricting the maximum-error audit to the driven window $t \geq 0.25$ s removed the extreme pre-actuation contact-stiffness spikes: heel share improved from 12/15 to 13/15, ground yaw moment from 13/15 to 15/15 and lumbar axial cut moment from 14/15 to 15/15. Horizontal ground force nevertheless remained 0/15, with its maximum η reduced from 131.86% to 23.98%. Thus the contact-force conclusion is not explained solely by a single initialization spike. Figure 8 summarizes the full-window maximum-error hierarchy across outputs.

Threshold sensitivity preserved the qualitative hierarchy. At 5% / 10% / 20%, the numbers of passing conditions were 10 / 15 / 15 for pelvis yaw, 15 / 15 / 15 for thorax yaw, 1 / 13 / 15 for free-side horizontal response, 0 / 12 / 13 for heel share, 0 / 14 / 15 for stance-axis angle, 0 / 0 / 10 for horizontal ground force, 13 / 13 / 14 for ground yaw moment, and 13 / 14 / 15 for axial lumbar cut moment.




Figure 8. Locked-normalized yaw-only error over 15 operating conditions. Bars show the median and maximum error ratio η by output. The 10% reporting criterion is output-specific; contact-force-related quantities show the largest failures.

3.4. Robustness and sensitivity

The reflected left-stance simulations provided an implementation check: yaw quantities changed sign as expected while scalar support measures agreed to numerical precision. The higher-drive condition (20 N·m ankle, -1.5 N·m hip yaw) also preserved the qualitative reduction result: final free-side response was -1.87 mm for locked, 11.13 mm for yaw-only and 10.17 mm for three-coordinate; final thorax yaw was -0.0299, 0.0153 and 0.0170 rad, respectively. The negative locked-model endpoint is phase-dependent rather than a persistently reversed response: over the 1.5 s transient, locked free-side x ranged from -12.64 to +16.28 mm and crossed zero before ending at -1.87 mm. This extends the finding beyond the nominal amplitude without turning the synthetic task into human validation.

The lumbar-impedance sweep makes the stabilization boundary explicit. At the nominal condition, relaxing transverse stiffness from 1000 to 100 N·m/rad caused the yaw-only approximation to exceed the 10% criterion for every audited output. At 100 N·m/rad, η was 30.37% for pelvis yaw, 17.23% for thorax yaw, 45.84% for free-side displacement, 32.97% for heel share, 79.25% for stance axis, 33.18% for horizontal ground force, 26.05% for ground yaw moment and 18.63% for axial lumbar cut moment. At 300 N·m/rad only thorax yaw and axial lumbar cut moment met 10%; at 1000 N·m/rad all listed outputs except horizontal ground force met 10%. The full-window 3000 N·m/rad case shows non-monotonic contact maxima caused by an early coefficient-change transient; over $t \geq 0.25$ s, heel share, horizontal ground force and ground yaw moment were 4.82%, 13.02% and 2.81%, respectively. The reduction result is therefore not a generic property of a yaw-only joint: it depends materially on effective transverse stabilization, and contact outputs must be distinguished from rotational kinematics. Figure 9 shows the corresponding kinematic trends, and Table 10 reports the full eight-output adequacy boundary across transverse-stiffness levels.

Sensitivity of yaw-only reduction to transverse lumbar stiffness




Figure 9. Sensitivity of the yaw-only reduction to transverse lumbar stiffness. (a) Maximum absolute yaw-only versus three-coordinate differences in pelvis, lumbar and thorax yaw; (b) maximum free-side horizontal-response difference. Values are evaluated on the common time base for transverse stiffness values of 100, 300, 1000 and 3000 N·m/rad. Axial stiffness and damping, and transverse damping, were evaluated separately in the full archived sweep. The response boundary is output-specific rather than a single monotonic stiffness threshold.

Table 10. Nominal-condition locked-normalized maximum-error ratio η [%] across transverse lumbar stiffness. Values $\leq 10\%$ satisfy the primary reporting criterion. The 100 N·m/rad column provides the full-output low-stiffness boundary requested by the sensitivity analysis.

Output	100	300	1000	3000
Pelvis yaw	30.37	13.87	4.94	2.37
Thorax yaw	17.23	8.96	2.79	0.78
Free-side horizontal response	45.84	24.46	7.41	2.28
Heel share	32.97	17.71	7.11	23.59
Stance-axis angle	79.25	17.46	7.41	5.11
Horizontal ground force	33.18	30.21	18.20	80.93
Ground yaw moment	26.05	11.63	3.96	13.27
Lumbar axial cut moment	18.63	7.14	2.55	2.86

The supplementary upper-body formulation passed its prescribed-yaw known-answer check: added mass was 12.615 kg, yaw inertia about the attachment frame was 0.343148 kg·m², and measured torque matched the analytical 0.171574 N·m value to numerical precision. All 45 operating runs completed the runtime, force-balance and energy-ledger gates. Over the full interval, yaw-only met the 10% criterion in 15/15 conditions for thorax yaw and 13/15 for pelvis yaw; the two pelvis exceedances were 11.19% and 11.95%. Support/free-side transferability narrowed substantially (Table 11). The actuation-window audit showed which failures were robust to removal of the initial settling period: pelvis yaw remained 13/15, thorax

yaw 15/15, free-side response 2/15, heel share 0/15, stance axis 0/15 and horizontal ground force 0/15, whereas ground yaw moment improved from 4/15 to 12/15 and lumbar axial cut moment from 14/15 to 15/15. Under η_{RMS} in the driven window, free-side response met the 10% RMS criterion in 14/15 conditions and ground yaw moment/lumbar moment in 15/15, while heel share remained 0/15 and stance axis 1/15. Thus some force/moment maxima were transient-sensitive, whereas the loss of heel-share and stance-axis transferability was persistent.

Table 11. Output-specific yaw-only performance under the 10% criterion after restoring rigid head and upper-limb inertia across the same 15-condition operating sweep. Compare with Table 9 for the primary campaign. The last column reports the median and maximum $E_B = \max_t |y_{\text{yaw}} - y_{3D}|$ in native units.

Output	Median η [%]	Maximum η [%]	Conditions $\eta \leq 10\%$	Median / max absolute deviation
Pelvis yaw	6.82	11.95	13 / 15	0.00198 / 0.00350 rad
Thorax yaw	1.30	8.94	15 / 15	0.000534 / 0.00317 rad
Free-side horizontal response	11.53	15.22	2 / 15	1.549 / 1.996 mm
Heel share	23.25	99.50	0 / 15	0.01061 / 0.09599
Stance-axis angle	16.33	24.83	0 / 15	0.000714 / 0.001019 rad
Horizontal ground force	100.06	142.75	0 / 15	2.988 / 15.692 N
Ground yaw moment	10.91	40.99	4 / 15	0.1592 / 0.7340 N·m
Lumbar axial cut moment	4.92	12.74	14 / 15	0.0578 / 0.2537 N·m

4. Discussion

4.1. What the comparison can establish

The study is a controlled model-reduction experiment, not a validation of a human lumbar spine. Its central contribution is a reproducible validity-boundary procedure: hold mass properties and task inputs fixed, verify the implementation independently, quantify reduction error by output, and then stress the decision across operating, parameter and inertial changes. This follows the intended-use logic emphasized in computational-biomechanics verification and validation guidance (Henninger et al. 2010; Lund et al. 2012; Hicks et al. 2015) and makes explicit a point also seen in other musculoskeletal-model comparisons: additional degrees of freedom can matter strongly for some outputs and tasks while adding little for others (Valente et al. 2015).

The present results sharpen that principle for lumbar reduction. Under nominal effective transverse stabilization, yaw-only closely reproduced pelvis/thorax yaw while horizontal shear remained architecture-sensitive. When transverse stiffness was reduced to 100 N·m/rad, even pelvis and thorax yaw failed the 10% criterion; when rigid upper-body inertia was restored, thorax yaw remained robust but several support/free-side outputs did not. The scientifically

useful result is therefore not that 'yaw-only is sufficient,' but the mapped set of conditions under which it is and is not sufficient.

If a mobile architecture lowers one moment exposure but increases another peak, the result is redistribution rather than uniform load reduction. If its effect is negligible for a defined task, a rigid model may be a defensible and less costly approximation. Either outcome is useful when accompanied by a reproducible validity boundary.

The nominal and robustness campaigns support a conditional reduction result. In the primary 15-condition sweep, yaw-only lumbar mobility met the locked-normalized 10% criterion for pelvis and thorax yaw in all 15 conditions, while horizontal ground force did not. The supplementary rigid upper-body inertia sweep retained 15/15 thorax-yaw passes under the 10% criterion but reduced the pelvis-yaw pass count to 13/15 and substantially narrowed transferability for free-side and support/contact outputs. The two pelvis-yaw exceedances were modest and both remained below 12%. Resultant lumbosacral moment exposure changed little in the nominal primary comparison even though axial moment organization changed substantially. The reduction decision therefore depends on the intended observable and on the inertial scope of the model.

4.2. Separation from contact and control assumptions

Contact modeling is a major boundary of the conclusion. In the primary sweep, the largest full-window heel-share, horizontal-force and ground-yaw errors occurred in the extreme contact-stiffness cases. The driven-window audit showed that some of those extreme ratios came from early penalty-contact settling, but horizontal ground force still failed 10% in all 15 conditions after $t = 0.25$ s. This persistent sensitivity is mechanically plausible because the regularized Coulomb law maps small differences in tangential foot velocity and normal force directly into horizontal shear. Kinematic yaw can therefore remain close while a contact-force channel differs appreciably. The result should be read as a warning against transferring a reduction validated on segment angles directly to ground-force or plantar-load outputs.

The right/left reflection served as a numerical safeguard rather than a biological side-comparison. Sign symmetry was verified before interpreting architecture effects. The higher-drive sweep likewise tested amplitude robustness, but the applied loads remain synthetic and should not be read as measured human joint moments.

4.3. Biological scope and limitations

An effective three-coordinate lumbar joint does not predict segmental vertebral loads. It uses independent linear springs and does not represent the coupled motion of the lumbar spine, in which axial rotation is accompanied by lateral bending and, to a lesser extent, flexion (Shin et al. 2013), nor the nonlinear load–displacement behaviour of the motion segment, with a low-stiffness neutral zone followed by stiffening towards end-range (Panjabi 1992b). It also concentrates all lumbar mobility at a single effective joint, whereas axial rotation is distributed unevenly across levels and limited by the orientation of the facet joints. These simplifications are acceptable for the small amplitudes simulated here but would need to be revisited before the protocol is applied to larger rotations.

The transverse-stiffness audit is the strongest limitation on the nominal yaw-only result. The three-coordinate reference is comparatively stiff in flexion and lateral bending at the nominal 1000 N·m/rad, which necessarily promotes agreement with a yaw-only architecture. The expanded audit makes this assumption testable rather than implicit: at 100 N·m/rad all eight audited outputs in the nominal condition exceed the 10% criterion, and pelvis/thorax yaw themselves rise to 30.37% and 17.23%. Yaw-only adequacy should therefore be claimed only for effectively stabilized transverse motion in the tested range, not for a freely compliant passive lumbar joint.

Inertial scope is a second boundary. The baseline represents 62.79 kg and omits explicit head, arms and an articulated neck. The supplementary sensitivity rigidly adds 12.615 kg of de Leva-based head/upper-limb mass and 0.343148 kg·m² of yaw inertia, producing 75.405 kg because the baseline pelvis/thorax partition is synthetic. These bodies are rigidly attached to the thorax, so the experiment tests inertial sensitivity rather than physiological shoulder/arm motion. The preserved baseline contact geometry also creates a short settling transient; the driven-window audit shows that this transient explains some contact-moment maxima but not the persistent free-side, heel-share and stance-axis sensitivity. Absolute magnitudes should therefore not be transferred directly to whole-body human turning. The supplementary campaign was intentionally not re-equilibrated by retuning passive rest angles or contact penetration, because doing so would combine the inertial perturbation with a change in support preload. The full-window results therefore include a short settling response, and persistent inertial effects are interpreted primarily from the parallel $t \geq 0.25$ s audit; early settling peaks are not presented as steady human-support behavior.

The reference three-coordinate joint is itself an unvalidated reduced spine: it concentrates lumbar motion at one effective joint, uses independent linear coordinates and does not reproduce vertebral coupling, neutral-zone nonlinearity, muscle recruitment or calibrated plantar tissue. The tested rotations are small to moderate (below 2° in the primary campaign; up to 6.31° pelvis and 7.47° thorax yaw in the upper-body sensitivity), which limits but does not remove those concerns. The next validation step is a matched human experiment using motion capture and force plates during small single-leg trunk rotations, with pelvis/thorax kinematics and ground reactions designated in advance as validation quantities.

4.4. Relevance for applied musculoskeletal modelling

For applied musculoskeletal modeling, the practical implication is output-specific. A reduced lumbar architecture may be defensible for pelvis/thorax orientation in a small, effectively stabilized rotational task yet unsuitable for horizontal ground force, plantar-load distribution or balance-related outputs under the same inputs. This is consistent with evidence that joint-model complexity should be selected according to task and intended prediction rather than maximized by default (Valente et al. 2015). The protocol provides a way to document that decision quantitatively before a reduced model is reused for another question.

No result here supports inference about individual muscle force, vertebral contact stress, low-back-pain mechanisms or injury risk; those questions require different model structure and experimental validation.

5. Conclusions

This study establishes an output-specific validity boundary for lumbar model reduction in a support-driven multibody task. Under the nominal 1000 N·m/rad transverse-stabilization surrogate, yaw-only mobility met the 10% maximum-error criterion for pelvis and thorax yaw in all 15 primary operating conditions, while horizontal ground force failed in all 15. That result is not universal: at 100 N·m/rad transverse stiffness every audited output in the nominal condition exceeded 10%, including pelvis yaw (30.37%) and thorax yaw (17.23%). Restoring rigid head and upper-limb inertia met the same 10% criterion for thorax yaw in 15/15 conditions and pelvis yaw in 13/15 but substantially reduced transferability of several support outputs; removing the pre-actuation settling interval did not restore free-side, heel-share or stance-axis adequacy. Maximum-error and RMS audits further show that some failures are brief transients whereas others persist over the trajectory. The broader contribution is therefore a reproducible decision procedure—mass-equivalent ablation, independent benchmark, reflection check, numerical refinement, parameter/inertial sensitivity and output-wise error metrics—for determining when a multibody joint can be simplified for a specified modeling question. Human validity remains a separate experimental problem.

Declarations

Data availability statement: The Article 2 dataset revision is publicly archived at <https://doi.org/10.5281/zenodo.23175126> (concept DOI 10.5281/zenodo.22957143), and the ModelicaHumanBodyPArts v0.10.0 software release is archived at <https://doi.org/10.5281/zenodo.22941517>. The published dataset contains ModelicaHumanBodyPArts_Article2_v010_REPRO_REV2_20261006.zip (SHA-256 425b117ba25de9bb1d3bb0517eb1b3e54770157f48ad0cd64119957050655255), including the preserved primary campaign, the supplementary upper-body source and 45/45 verified upper-body operating runs, reproduction tests/runners, logs and machine-readable summaries. The RMS, actuation-window and expanded transverse-stiffness results reported in this manuscript are deterministic post-hoc analyses of the frozen trajectories; they do not alter model source or raw simulation outputs. Research data and non-software documentation are CC BY 4.0; software/code remains MIT licensed. The frozen reproduction commit is <https://github.com/antontofic000-creator/ModelicaHumanBodyPArts/tree/cdf17e684e4f0c79277651d88bf9c7bac397243d/studies/article2>; the maintained study path is <https://github.com/antontofic000-creator/ModelicaHumanBodyPArts/tree/main/studies/article2>.

Preprint availability: A public preprint version of this study is available at <https://doi.org/10.5281/zenodo.23185698> (concept DOI 10.5281/zenodo.22958895).

Funding: No external funding was received for conducting this study or preparing the manuscript.

Disclosure statement: T.A. El Halabi is the originator of THE LYNKORR PROJECT, developer of ModelicaHumanBodyPArts, and author of a commercially available book describing the Lynkorr framework. These relationships are disclosed as potential non-financial

and commercial competing interests. N. Chahine declares no competing interests. The authors report no external or commercial funding for this study.

Author contributions (CRediT): Toufic Antoun El Halabi: Conceptualization; Methodology; Software; Validation; Formal analysis; Investigation; Data curation; Visualization; Writing – original draft; Writing – review & editing; Project administration. Nathalie Chahine: Methodology (physiological parameterisation); Validation (physiological plausibility of outputs); Writing – original draft (physiological scope and applied-relevance sections); Writing – review & editing.

Ethics approval and consent: This is a computational study using synthetic/source-derived parameters and no new human participant data or animal experiments. Ethics approval, informed consent and consent for publication are not applicable.

Declaration of generative AI use: ChatGPT (OpenAI, GPT-5.6 Sol) was used by T.A. El Halabi for manuscript organization, language editing, literature-search support and assistance with plotting/reproduction code. Claude (Anthropic) was used by N. Chahine for review structuring and language editing. Neither system was treated as an author or as an independent source of scientific evidence. All references, numerical results, analyses, code and scientific claims were checked by the authors. Research-result figures are deterministic programmatic diagrams or plots derived from author-developed model definitions and archived data; no generative-AI image model was used to create or alter research-result figures.

References

- Beaucage-Gauvreau E, Robertson WSP, Brandon SCE, Fraser R, Freeman BJC, Graham RB, Thewlis D, Jones CF. 2019. Validation of an OpenSim full-body model with detailed lumbar spine for estimating lower lumbar spine loads during symmetric and asymmetric lifting tasks. *Comput Methods Biomech Biomed Engin.* 22(5):451–464. doi:10.1080/10255842.2018.1564819.
- Bova M, Massaro M, Petrone N. 2020. A three-dimensional parametric biomechanical rider model for multibody applications. *Appl Sci.* 10(13):4509. doi:10.3390/app10134509.
- Bruno AG, Bouxsein ML, Anderson DE. 2015. Development and validation of a musculoskeletal model of the fully articulated thoracolumbar spine and rib cage. *J Biomech Eng.* 137(8):081003. doi:10.1115/1.4030408.
- Caron S, Pham Q-C, Nakamura Y. 2015. Stability of surface contacts for humanoid robots: closed-form formulae of the contact wrench cone for rectangular support areas. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). Seattle (WA): IEEE. p. 5107–5112. doi:10.1109/ICRA.2015.7139910.
- Christophy M, Faruk Senan NA, Lotz JC, O'Reilly OM. 2012. A musculoskeletal model for the lumbar spine. *Biomech Model Mechanobiol.* 11:19–34. doi:10.1007/s10237-011-0290-6.
- de Leva P. 1996. Adjustments to Zatsiorsky–Seluyanov's segment inertia parameters. *J Biomech.* 29(9):1223–1230. doi:10.1016/0021-9290(95)00178-6.
- Delp SL, Anderson FC, Arnold AS, Loan P, Habib A, John CT, Guendelman E, Thelen DG. 2007. OpenSim: open-source software to create and analyze dynamic simulations of movement. *IEEE Trans Biomed Eng.* 54(11):1940–1950. doi:10.1109/TBME.2007.901024.

- de Zee M, Hansen L, Wong C, Rasmussen J, Simonsen EB. 2007. A generic detailed rigid-body lumbar spine model. *J Biomech.* 40:1219–1227. doi:10.1016/j.jbiomech.2006.05.030.
- Dukkipati ST, Driscoll M. 2026. An integrated rigid-flexible body dynamic approach to computationally efficient musculoskeletal modeling and muscle recruitment simulation of the lumbosacral spine and torso. *IEEE Trans Biomed Eng.* 73(5):1877–1889. doi:10.1109/TBME.2025.3617301.
- El Halabi TA. 2026. ModelicaHumanBodyPArts v0.10.0 [software]. Zenodo. doi:10.5281/zenodo.22941517.
- Henninger HB, Reese SP, Anderson AE, Weiss JA. 2010. Validation of computational models in biomechanics. *Proc Inst Mech Eng H.* 224(7):801–812. doi:10.1243/09544119JEIM649.
- Hicks JL, Uchida TK, Seth A, Rajagopal A, Delp SL. 2015. Is my model good enough? Best practices for verification and validation of musculoskeletal models and simulations of movement. *J Biomech Eng.* 137:020905. doi:10.1115/1.4029304.
- Lerchl T, Nispel K, Baum T, Bodden J, Senner V, Kirschke JS. 2023. Multibody models of the thoracolumbar spine: a review on applications, limitations, and challenges. *Bioengineering.* 10:202. doi:10.3390/bioengineering10020202.
- Lehner M, Eberhard P. 2006. On the use of moment-matching to build reduced order models in flexible multibody dynamics. *Multibody Syst Dyn.* 16(2):191–211. doi:10.1007/s11044-006-9018-2.
- Lund ME, de Zee M, Andersen MS, Rasmussen J. 2012. On validation of multibody musculoskeletal models. *Proc Inst Mech Eng H.* 226(2):82–94. doi:10.1177/0954411911431516.
- Otter M, Elmqvist H, Mattsson SE. 2003. The new Modelica MultiBody library. In: *Proceedings of the 3rd International Modelica Conference*. Linköping (Sweden): Modelica Association. p. 311–330.
- Panjabi MM. 1992a. The stabilizing system of the spine. Part I. Function, dysfunction, adaptation, and enhancement. *J Spinal Disord.* 5(4):383–389. doi:10.1097/00002517-199212000-00001.
- Panjabi MM. 1992b. The stabilizing system of the spine. Part II. Neutral zone and instability hypothesis. *J Spinal Disord.* 5(4):390–397. doi:10.1097/00002517-199212000-00002.
- Raabe ME, Chaudhari AMW. 2016. An investigation of jogging biomechanics using the full-body lumbar spine model: model development and validation. *J Biomech.* 49:1238–1243. doi:10.1016/j.jbiomech.2016.02.046.
- Seth A, Hicks JL, Uchida TK, Habib A, Dembia CL, Dunne JJ, Ong CF, DeMers MS, Rajagopal A, Millard M, et al. 2018. OpenSim: simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PLoS Comput Biol.* 14(7):e1006223. doi:10.1371/journal.pcbi.1006223.
- Shin J-H, Wang S, Yao Q, Wood KB, Li G. 2013. Investigation of coupled bending of the lumbar spine during dynamic axial rotation of the body. *Eur Spine J.* 22:2671–2677. doi:10.1007/s00586-013-2777-6.

- Valente G, Pitto L, Stagni R, Taddei F. 2015. Effect of lower-limb joint models on subject-specific musculoskeletal models and simulations of daily motor activities. *J Biomech.* 48(16):4198–4205. doi:10.1016/j.jbiomech.2015.09.042.
- Vierneisel M, Frie L, Eberhard P. 2026. Model Order Reduction - also with many inputs and outputs? *Multibody Syst Dyn.* 66:349–364. doi:10.1007/s11044-025-10088-8.
- Watt AA, Callaway AJ, Williams JM. 2022. In vivo through-range passive stiffness of the lumbar spine: a meta-analysis of measurements and methods. *Med Biol Eng Comput.* 60(8):2133–2157. doi:10.1007/s11517-022-02609-w.
- White AA, Panjabi MM. 1990. *Clinical biomechanics of the spine*. 2nd ed. Philadelphia (PA): J.B. Lippincott.