MOTION · ACTUATION

Technical Report TR-2026-41
Research / Modelled Results · Preprint v1.4
19 August 2026

MOTION · ACTUATION ENERGETICS

Joules per Punch

Where the energy of a humanoid strike goes, why a segmented spine costs more than one waist joint at every demand both can meet, and why it is still the only body that throws the hardest strike.

The Charlot Lab · Institute for Physical AI @ John Bailey Institute

Takes the sum E_task = E_compute + E_actuation from the total-task-energy bench to a strike. A single-file bench, eight printed gates, no dependencies. Every bench and every run receipt behind the figures below is released into the public domain: huggingface.co/datasets/physicalai-bmi/punch-energetics.

serial joints share motion, not torque

copper multiplies as N to the 8/3

joint count is free once actuation spans joints

the reward, not the policy

a feasibility technology before an efficiency one

In January 2026 two Unitree G1 humanoids boxed at CES, and the commentary settled on one diagnosis: the motion looks wrong because the control policy is weak [7]. The kinematic diagnosis runs deeper. The classical force decomposition of the boxer's straight punch places 38.5 percent of its force in leg drive and 37.3 percent in trunk rotation, against 24.1 percent in the arm [1], and the base G1 carries one yaw joint where that trunk contribution lives [6]. The obvious engineering response is a segmented spine. This report prices that proposal against an identified electrical power model and finds it inverted: on a fixed strike, defined by Khatib's operational-space effective mass so that a fast, empty arm cannot score [2], **every segmented trunk tested costs more than the single waist joint**, from +13 percent at two vertebrae to +4,740 percent at seventeen under a fixed actuator-mass budget. The mechanism is that serial joints each transmit the full distal torque, so N rotary modules pay N copper terms while each smaller module's motor constant worsens as mass to the five-thirds; the two compound to roughly N to the eight-thirds. The conclusion survives sweeping the scaling exponent to its most generous physical value, where an eight-segment spine still pays +219

percent. What does move the bill is actuation topology: with few multiarticular actuators spanning many joints, seventeen vertebrae cost within a few joules of eight, because the copper bill follows actuator count rather than joint count. The demand frontier then reverses the verdict in kind rather than degree: sweeping the required momentum upward with the speed band fixed, every single-joint and eight-joint body tested stops delivering at 18 kg·m/s, and the seventeen-segment multiarticular trunk is the only architecture that delivers 22, with replicated witnesses. The segmented trunk is a feasibility technology before it is an efficiency one. Two further results frame the finding. Scored on momentum alone, the energy-optimal strike for this body leaves the torso nearly still, so the flailing seen in the ring is close to optimal play for the objective in use. And the compute term of the strike is 1.5 to 2.5 percent of the total, with the arithmetic itself six orders of magnitude below platform idle, so policy compression optimises a term that is not there.

1. What question does this report answer?

The robot boxing matches of 2025 and 2026 put a humanoid's whole-body coordination on public display, and the display reads as flailing: fast hands, still torso, feet that exist to not fall over [7]. The standing explanation blames the policy. A better explanation starts from the body. Filimonov's force decomposition on master boxers assigns roughly three-quarters of straight-punch force to leg drive and trunk rotation together [1], and the base G1 offers a single waist yaw joint on that entire power path, with the EDU variant adding roll and pitch [6]. A policy cannot express a motor program its body has no channel for.

From there the engineering proposal writes itself: give the robot a spine. Replace the single waist joint with a stack of vertebral modules and let the trunk deliver its 37 percent. This report asks whether that proposal survives contact with an actuator, in the only currency a battery understands. It prices the same strike on four architectures and reports where every joule goes.

2. How is a punch priced so that bodies can be compared?

Fist kinetic energy will not serve as the task, because an arm-only punch can be fast and still weak. What a target receives is impulse, and impulse is set by the effective mass the body presents at the contact point along the punch direction. Khatib's operational-space formulation gives that quantity exactly [2]:

Quantity	Definition	Value used
Effective mass	m_eff = 1 / (uᵀ J M(q)⁻¹ Jᵀ u), the inertia presented at the fist along direction u	floor of 2.4 kg, 7 percent of the 35 kg body
Delivered momentum	P = m_eff · v_fist	P* = 9.0 kg·m/s, identical for every architecture
Fist speed band	a strike, not a shove and not a slap	2.8 to 3.8 m/s at contact, 280 ms from guard

Each constraint closes a loophole the optimiser otherwise takes. Momentum alone lets a body trade effective mass against speed one for one, and the search returns either a fast, nearly massless arm or a slow shove with the whole torso behind it; the speed band excludes both. The effective-mass floor is what a coach means by putting your weight behind it, and the sport-science literature grounds the choice: trained strikers present roughly twice the effective mass of novices at similar hand speeds [9]. Every architecture must deliver the same strike; whichever spends fewer joules wins.

Every architecture is also granted a perfect energy-optimal controller, found by a multi-start constrained search over strike profiles, so the comparison is the floor the body imposes rather than any policy's skill. The electrical cost per actuator follows the identified G1 arm power model: copper loss as torque squared over the motor constant, viscous and Coulomb friction, positive mechanical work, and idle [3]. Loss magnitudes are calibrated so a nominal reach lands near the 71.5 joules measured on the physical G1 [4], and the sibling bench carries that calibration [11]. Absolute joules are indicative; the comparisons between architectures are the result.

3. What does the baseline strike cost, and where does it go?

Architecture A is the G1 as built: a rigid trunk on one waist joint driven by one 1.8 kg motor. Its energy-optimal strike delivers the target momentum at 3.19 kg of effective mass and 2.83 m/s, and the bill decomposes as follows.

Term	Joules	Share	What it is
Positive mechanical work	106.4	39%	the kinetic energy of the strike itself, plus work against coupling
Copper loss	83.0	31%	torque squared over the motor constant, of which the trunk motor pays 7.3 J
Bearing friction	52.4	19%	viscous and Coulomb, per joint
Motor viscous	23.3	9%	speed-squared loss referred to each actuator's output
Idle	5.6	2%	drive electronics and holding current over the strike window
Total actuation	270.7	100%	the number every other architecture is measured against

Two features of this decomposition drive everything that follows. Copper is the second-largest term and the only one that depends on how torque is produced rather than on the motion itself. And the trunk's share of it is small in the baseline, which is the first hint of Section 4's result: even a perfect controller barely uses the one trunk joint it has.

4. What happens when momentum is all that is scored?

Drop the effective-mass floor, the way a reward function scores a hit sensor or a momentum estimate, and re-run the search. The optimum barely changes: 3.22 kg at 2.80 m/s, 268.3 joules, and the trunk motor pays 6.9 of the 82.0 joules of copper, 8 percent of the copper bill for the joint that anatomy says should carry 37 percent of the force.

Driving twelve kilograms of torso is not worth the joules when a lighter, faster arm meets the same score. The flailing in the ring is therefore close to optimal play for the objective in use, and this is a result about morphology and scoring before it is a result about any policy. Training longer against the same reward converges to the same arm punch. The correction belongs in the objective: score the effective mass presented at contact, a quantity computable from the model the controller already carries [2], and the trunk becomes worth paying for.

This revision ran the learner, and the objective needs more precision than that sentence gave it [14]. Cross-entropy policy search over the same strike, thirteen reward conditions, five seeds each, six of thirteen registered predictions falsified. Scored on momentum alone the learner converges to a 7.5 m/s whip with 0.67 kg behind it, a stronger form of the claim above. Naming the effective mass then fails in either polarity: as an additive bonus it buys a heavy, slow shove at every weighting tried, and as a penalty below the floor it buys the same shove, because above the floor speed is unpriced and heavy-and-slow satisfies a constraint cheaply. An optimiser under a hard constraint minimises energy at the task; a learner maximises reward, and a penalty-satisfied shove is reward-cheap. What delivers the strike is the multiplicative form, momentum times a weight-behind-it gate, arriving as a curriculum and searched with exploration that outlives it: four of five seeds deliver a median 13.1 kg·m/s at 2.39 kg and 5.46 m/s, past the target this report set. Remove the curriculum and the gate starves at stillness; remove the exploration and the tightening gate outruns the collapsing search. The learned equilibrium parks at the gate's ramp top, 2.39 against 2.40 kg, a reminder that a success threshold placed at the reward's own boundary measures the reward rather than the behaviour.

5. What does adding vertebrae cost?

Architecture B is the naive spine: N trunk joints, each driven by its own rotary module, the arrangement every modular-actuator catalogue suggests. The actuator mass budget matters, so both conventions are run. B splits the baseline's 1.8 kg of trunk motor across the N modules; B' gives every module the full 1.8 kg and lets trunk mass grow.

Trunk joints	B: fixed total motor mass	vs A	B′: full-size modules	vs A
1 (baseline A)	270.7 J		270.7 J	

2	306.1 J	+13%	277.2 J	+2%
3	394.2 J	+46%	286.5 J	+6%
5	749.8 J	+177%	303.8 J	+12%
8	1,941.8 J	+617%	317.3 J	+17%
12	5,325.8 J	+1,867%	359.1 J	+33%
17	13,101.9 J	+4,740%	405.0 J	+50%

Seventeen is the human thoracolumbar count, and at seventeen the split-budget spine pays 48 times the baseline for the identical strike. The full-size-module column avoids the catastrophe by spending 28.8 kg of additional motor mass, 30.6 kg of trunk motor in all, to reach a 50 percent penalty. There is no crossover to find: the segmented trunk never dips below the single joint at any count, under either convention.

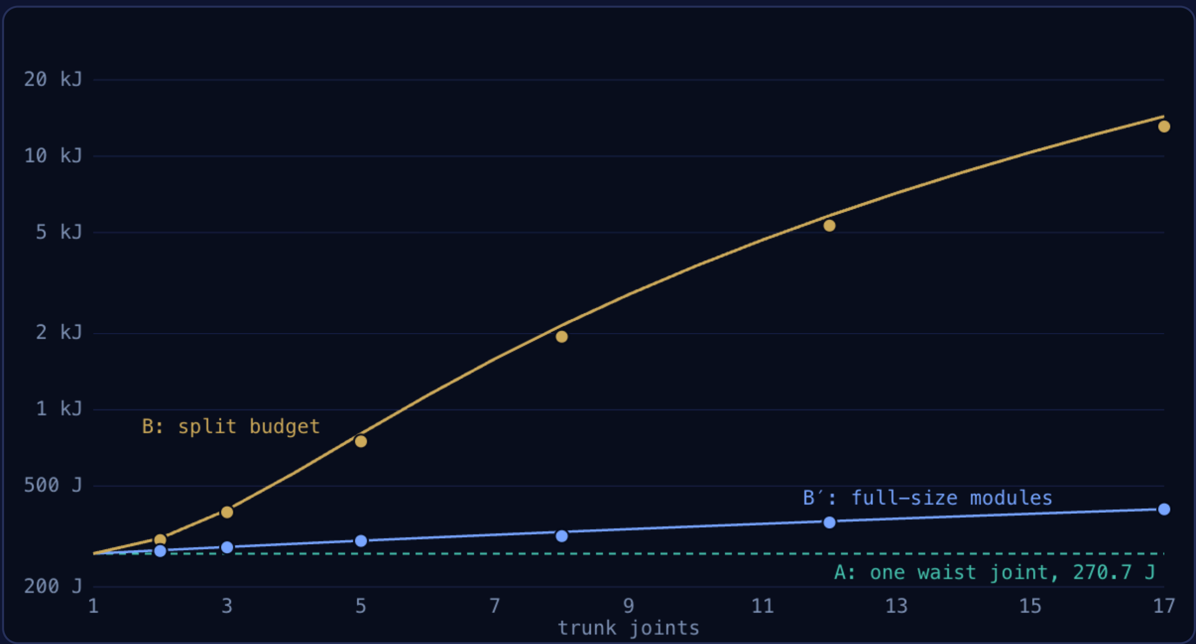

motor scaling exponent **1.67**

At exponent 1.67, eight split-budget vertebrae price at 2153 J, 8.0x the single joint. The dots are the full bench at 1.67.

Strike energy against vertebra count, log scale. The curves are the closed-form copper substitution of Section 6 anchored on the measured baseline decomposition; the dots are the full bench at exponent five-thirds, optimiser and all, which lands at or under the closed form because the search reallocates the strike away from the expensive trunk. Only the split-budget curve responds to the slider: full-size modules keep the baseline motor constant at every count.

6. Is the motor scaling law doing the work?

The result rests on two pieces of physics, so each is stated and then stressed.

Serial joints share motion, not torque. In a serial chain every joint transmits the full torque of everything distal to it. Splitting one waist joint into N vertebrae does not divide the load; it replicates it, so N modules pay N copper terms of the same size where the baseline paid one. Bearing friction moves the other way, since the same total rotation split N ways costs 1/N in a speed-squared term, and the bench shows exactly that in its bearing column. Copper dominates.

Smaller motors are worse at torque. For geometrically similar machines the motor constant squared, torque squared per watt of copper, scales as actuator mass to the five-thirds [8]. Under a fixed mass budget each of N modules is 1/N the size, so each pays a further factor of N to the five-thirds. The two effects compound: the trunk copper multiplies as N to the one plus the exponent, roughly N to the eight-thirds at the geometric-similarity value.

The exponent is the load-bearing assumption, so it is swept across its physical range. An exponent of 1.00 says copper splits perfectly with mass, crediting the segmented design with everything geometry refuses it.

Trunk joints	exponent 1.00	1.33	1.67 (geometric similarity)	2.00
3, split budget	+30%	+36%	+43%	+52%
8, split budget	+219%	+365%	+604%	+997%

The geometric-similarity column re-prices the architecture inside the sensitivity harness with its own multi-start search, so its +604 percent for eight joints sits a few percent under the main table's +617 for the same body: that spread is the search scatter, printed rather than reconciled. At the most generous exponent physically available, an eight-segment spine still pays more than three times the baseline. The N-fold replication of the torque bill carries the result on its own, and that term is set by the topology of a serial chain rather than by any motor technology. Better magnets do not repeal it.

7. Which winding actually runs out first?

Every figure so far is a joule or a sum of watts, and a sum is the wrong shape for a thermal question. Copper loss is heat deposited in a particular winding, and a motor stops when *one* of its windings reaches its insulation limit, not when the machine's total reaches anything. A whole-body peak of 2.6 kilowatts spread across twenty actuators is a different machine from the same 2.6 kilowatts in one, and until now this report could not tell them apart: it published a whole-body total and a trunk-group sum, and neither is a statement about any single motor.

The bench now reports the **peak electrical power of the hottest single actuator** in the strike, alongside the two sums. It is a maximum over actuators rather than a total across them, and it costs nothing to compute because the per-actuator terms were already being summed. **The result is a second, independent reading of this report's central finding, in heat rather than joules.**

The hottest winding is in the arm, and for most bodies it stays there. For the single-waist baseline it is an arm actuator at 617 watts, three and a half times the 184-watt trunk drive and nowhere near the 2,686-watt whole-body figure. Every architecture that keeps full-size trunk motors agrees: the fixed-per-module spines hold between 620 and 655 watts from two segments to seventeen, and every multiarticular arrangement holds between 629 and 651. Adding vertebrae to those bodies does not change which motor is closest to its limit.

Only the split-mass family moves, and it moves exactly where the joules do.

trunk segments, fixed total actuator mass	energy vs A	hottest winding	which actuator
1 (the baseline)	–	617 W	arm
2 (0.90 kg/module)	+13%	653 W	arm
3 (0.60 kg/module)	+46%	666 W	**trunk**

trunk segments, fixed total actuator mass	energy vs A	hottest winding	which actuator
5 (0.36 kg/module)	+177%	1,426 W	trunk
8 (0.23 kg/module)	+617%	3,512 W	trunk
12 (0.15 kg/module)	+1,867%	6,211 W	trunk
17 (0.11 kg/module)	+4,740%	**11,260 W**	trunk

The crossover is at three segments. Below it the machine is arm-limited and the spine is thermally irrelevant; at and above it the spine becomes the part closest to failure, and then gets worse by a factor of seventeen while the energy bill grows by a factor of thirty. **The segmentation penalty has a thermal signature, and it is a change in which motor is about to stop.** That is not a new claim, it is Section 5's claim arriving through a different quantity, which is the strongest form of agreement a single bench can offer: the two readings share a dynamics core but not an argument.

It also sharpens what the multiarticular arrangement buys. Seventeen joints driven by three tendon actuators keep their hottest winding at 630 watts, statistically indistinguishable from the single-joint baseline's 617. **Vertebrae are free thermally for the same reason they are free energetically**, and the reason is the same one: the bill follows actuators, and that family does not add any.

What this does not license. The hottest winding is a measured quantity of this model; the *temperature* it reaches is not, because that needs a thermal model of the specific motor, and the arm actuators here are a third the mass of the trunk modules and would have their own constants. This report states which winding is closest to its limit and by how much it leads the others. It does not state a time to that limit, and a companion bench that does so treats every such figure as modelled rather than measured. **The method that would settle it exists, and it is not ours.** Lin, Qian, Luo and Liang [16] identify a whole-body lumped-parameter thermal network for a quadruped by least-squares fitting against measured motor temperatures, reporting a mean squared error below 2.5 $°C^2$ across conditions including a payload, and a 2026 reinforcement-learning paper [17] drives that identified model into a locomotion policy. Two features of their model deserve more attention than their constants would: their motor nodes are thermally *coupled* to one another and to the onboard computer, and their convection term depends on body velocity. A bench that treats each winding as isolated and still is optimistic on the first count and pessimistic on the second, which is a reason to trust an ordering over a duration. This review did not locate that identification published for a humanoid, nor per-actuator thermal data for the machine these numbers are calibrated against.

8. What is a spine actually for?

Architecture C is the biological arrangement read as an engineering specification: many kinematic degrees of freedom driven by few actuators that each span multiple joints, with elastic elements in series. A tendon crossing eight vertebrae is one tension, one copper bill, applied along the whole span; the price is that one actuator can only command the block-mean torque of the joints it crosses, and the bench charges that residual against a 35 percent feasibility cap.

Configuration	Strike energy	vs A	Where the joules moved
8 joints, 2 actuators	284.1 J	+5%	copper back near baseline: 8.5 J on the trunk
12 joints, 3 actuators	295.5 J	+9%	a third actuator adds its own copper and idle
17 joints, 3 actuators	291.9 J	+8%	five more vertebrae than the row above, for a rounding error
17 joints, 4 actuators	338.8 J	+25%	the fourth actuator, not the seventeen joints, moves the bill

Once actuation spans joints, joint count stops costing anything. Eight joints on two actuators, twelve on three and seventeen on three land within four percent of one another, because the copper bill follows the actuator count. Vertebrae become free, and what they buy is travel, compliance and impact tolerance. This is the precise sense in which the human spine is not "more degrees of freedom": it is a continuum actuator, muscle sheets spanning many segments with elastic tissue in series. Its robot equivalent is a tendon, not a stack of rotary modules.

Series elasticity is the other lever, and it is priced by a quantity that is easy to get wrong. A linear element storing energy E through a winding travel of theta carries a peak torque of 2E over theta, so the copper cost of loading it falls quadratically with the travel available and does not fall with winding time: winding slowly at the same torque holds the same torque-squared loss for longer. Priced by travel, elastic pre-load on the baseline body reaches a 1 percent saving at a 75 percent elastic share, real but small at this strike energy. Priced by duration, the same model returns large fictitious savings, which is the version of this proposal that does not survive its own arithmetic. Segmentation earns its keep exactly where it buys winding travel for such an element, which is a second, independent reason the useful spine is tendon-driven.

A thirty-draw replication at this section's operating point completes the elastic story from both directions [15]. On the multiarticular trunk, whose actuators genuinely work, removing the elastic element costs 24 joules: 295.0 J with no elastic against 271.2 J at a 75 percent share, with trunk copper falling from 42.1 to 2.6 J. On the rigid single joint the same element always costs about 6 joules, because its trunk barely moves and the wind-up is pure overhead. The spring pays on the drive that works and taxes the one that idles. The same replication is a discipline result: a fuller search moved every architecture's optimum 2 to 3 percent below its published single-draw value, orderings unchanged, so single-draw gaps under five percent are read as ties throughout.

The demand frontier is where the segmented trunk finally earns its keep [15]. Sweeping the required momentum upward with the speed band fixed, so the effective-mass floor climbs from 2.4 toward 7.1 kg, every single-joint and eight-joint body tested finds its last witness at 18 kg·m/s; at 22, across eight or more independent search attempts each, none delivers. The seventeen-segment trunk on three actuators delivers 22 in three of six draws, at 419 joules and 75 degrees of trunk travel, and at matched search budget it does so while costing 5 percent more than the single joint at low demand. The spine's case was never the cheap strike. It is the strike no other body can throw.

Peak power is where the element earns its place, and this revision adds the measurement. At a 75 percent elastic share the trunk actuator's peak electrical draw falls from 184 to 29 watts, a 6.4-fold reduction, while the wind-up relocates a mean of 8 watts into the slow phase. The whole-body peak stays near 2.6 kilowatts throughout, because this strike lives in the arm. The spring resizes the drive it crosses; it does not resize the strike. That is the precise form of the stretch-shorten claim: removal of peak power at the actuator the element spans, which is the component a motor drive is sized by, with a small mean relocated to a phase that has time to spare.

9. Where does compute sit in a strike?

The sum this series prices is E_task = E_compute + E_actuation, and the compute term of a punch is small in a specific, quantifiable way. A 23-degree-of-freedom locomotion-class policy at 50 Hz executes about 175,000 multiply-accumulates per step. Over the 280 ms strike:

Platform duty	Platform energy	Arithmetic energy	Share of E_task	A 10x smaller policy saves
Orin NX class, policy only, 15 W	4.20 J	2.5×10^{-6} J	1.5%	under a millionth of a percent
with perception, 25 W	7.00 J	2.5×10^{-6} J	2.5%	under a millionth of a percent

The arithmetic sits six orders of magnitude below the platform's own idle, and the platform itself is 1.5 to 2.5 percent of the strike. Compressing the policy therefore optimises a term that is not there, which reproduces on a strike what the embodied-efficiency literature reports for manipulation [5] and what the sibling bench measures per completed reach [11]. The compute lever that does exist is duty cycle: platform watts times seconds is the whole compute bill, so gating the platform moves it and shrinking the network does not. TR-2026-40 states the same bound from the other side, as the reciprocal of the non-compute share [13].

10. What did the bench have to survive?

A dynamics model that is wrong produces confident nonsense, so the bench refuses to print results until eight gates pass, and the run output opens with them. Two of the eight exist because an earlier revision of the bench produced a finding that was fake, and both failures generalise beyond this bench.

Gate	What it pins	Result
Energy conservation	free rotation about a vertical axis holds kinetic energy over 80 ms	drift 3.2×10^{-6}
Inverse dynamics identity	recursive Newton–Euler equals M(q) times acceleration at zero velocity	2.4×10^{-16}
Work–energy along the strike	net mechanical work equals the kinetic energy created, on the actual profile	6.4×10^{-4}
Analytic effective mass	a single link presents I over r squared, tangentially	exact
Locked–joint control	a 5–joint trunk with 4 joints frozen reproduces the 1–joint chain	1.8×10^{-16}
Degenerate–direction guard	a near–motionless fist is rejected as inadmissible	rejects at 2.8×10^{-4} m/s
Scaling law wiring	halving module mass costs 3.17x in copper, as the exponent demands	exact
Optimiser soundness	tightening a constraint never lowers the reported optimum	slack −0.6%

The degenerate-direction guard. The first revision reported a 98 percent energy saving for a three-segment spine. Near a singular punch direction the quadratic form under Khatib's effective mass goes to zero, so the effective mass diverges, and the momentum target is met by a fist moving at a third of a millimetre per second. The optimiser found that attractor immediately, and the row it produced looked exactly like a result. The general rule: a physically unbounded quantity in an objective is an optimiser attractor, and it must be bounded explicitly rather than trusted to be unreachable.

Optimiser soundness. A later revision returned a strictly more constrained run 7 percent cheaper than a less constrained one, which no true optimum permits: tightening a constraint shrinks the feasible set.

The cause was a redundant radial dimension in the search parameterisation that the scale solver already handled, so the descent wandered among local minima and the wandering read as a morphology difference. The gate that now enforces monotonicity is a one-line check any constrained search can carry, and it converts a silent convergence failure into a loud one.

11. What is established, and what would extend it?

Established, within the model. That a fixed strike admits a clean architecture comparison once effective mass, momentum and speed are all pinned (Section 2). That the segmented, per-joint-actuated trunk costs more than a single waist joint at every count tested, under both mass conventions, and that the penalty survives the scaling exponent swept to its most generous value (Sections 5 and 6). That the copper bill follows actuator count once actuators span joints, making joint count free (Section 8). That elastic storage is priced by winding travel rather than winding time (Section 8). That the compute share of a strike is 1.5 to 2.5 percent with the arithmetic six orders below platform idle (Section 9). That elastic pre-load removes peak electrical power at the actuator it crosses, 184 to 29 watts at a 75 percent share, while the whole-body peak moves 3 percent (Section 8). That elastic storage saves 8 percent on the multiarticular trunk and costs on the rigid one, and that above 18 kg·m/s the seventeen-segment trunk is the only body tested that delivers the strike at all (Section 8). Replication discipline for all of it: cross-architecture gaps under five percent are search scatter, a thirty-draw replication moved every optimum down 2 to 3 percent with orderings unchanged, and feasibility evidence pools one-sidedly, a found strike standing forever while absence means only that the draws attempted found none.

Open. Whether a tendon-driven trunk's routing friction and residual-torque cost, charged here as a block-mean constraint with a 35 percent cap, hold at those values on hardware. Whether a gradient policy class and richer dynamics reproduce the reward-form result of Section 4, which cross-entropy search over this strike has now measured: additive and penalty rewards converge to a shove, and the gated curriculum delivers the strike [14]. And the anchor itself: this review did not locate a published joules-per-strike measurement for any humanoid, so the absolute scale here rests on the reach measurement of the identified power model's platform [4]. An instrumented strike, pack current against a force plate through contact, would replace a calibrated scale with a measured one and is within reach of any lab that owns the robot these matches made famous.

Prior art and position. Operational-space effective mass, serial-chain dynamics, motor scaling and series-elastic actuation are established fields. The punch force decomposition is forty years old [1], and tendon-driven musculoskeletal humanoids are a mature research line whose energetic rationale this result supports from the pricing side. What this report contributes is the joining: the strike stated as an effective-mass task, the segmented trunk priced per joule against its single-joint baseline under an identified loss model, the actuator-count reading of the multiarticular alternative, the per-actuator maximum that gives that reading a second and independent form in heat, and the two optimiser failures converted into reusable gates. This review did not locate these stated together in the literature it read.

12. Conclusions

A punch prices a body. Asked to deliver a fixed strike, a humanoid with one waist joint spends 271 joules, and every segmented spine bolted onto the same torso spends more, because a serial chain replicates torque down its whole length and torque squared is what copper charges for. The number that governs the bill is not how many joints the trunk has but how many actuators drive it, and the arrangement that makes vertebrae free, few tension elements spanning many joints with elastic series compliance, is the one the biological spine has used all along. The robot in the ring is not failing to use a spine it lacks; scored on momentum, it is correctly declining to pay for the torso it has. The cheapest improvements to the spectacle are therefore an objective that scores effective mass at contact, and a trunk actuated like a tendon rather than like a column of servos. The joules were never in the network. On a strike, as on a reach, they are in the windings, and the windings answer to topology. That last sentence turns out to be literal, and v1.4 makes it so: asked which single winding is closest to its limit rather than what the machine spends in total, the bench answers **an arm actuator, at 617 watts**, for every body that keeps full-size motors, including seventeen vertebrae driven by three tendons. Only the spine that splits a fixed actuator budget takes the lead away from the arm, at three segments, and by seventeen it is at 11,260 watts in one winding. The penalty for segmenting a trunk is not only a larger bill. It is a change in which motor stops first.

13. The forcing function

What is bounded	The physics that sets it	The engineering change that moves it	What becomes possible
A segmented, per-joint-driven trunk pays 13 to 4,740 percent over one waist joint	Serial joints each transmit the full distal torque, and copper charges torque squared per actuator, worsened by mass to the five-thirds under a split budget	Multiarticular actuation: one tension element spanning many joints, so the copper bill scales with actuator count	Vertebral count becomes free, and the frontier shows what it buys: above 18 kg·m/s of delivered momentum the seventeen-segment multiarticular trunk is the only body tested that delivers the strike at all, at a 7 percent premium at low demand
The hottest winding in the strike is an arm actuator at 617 W, and every body that keeps full-size motors is arm-limited, including seventeen vertebrae on three tendons	A motor stops when ONE winding reaches its insulation limit, so the binding quantity is a maximum over actuators and not the whole-body total. The strike's speed demand lives in the arm, so the arm's torque-	Cool the actuator the maximum is in rather than the one the mass is in; or raise the arm's motor constant at fixed mass. For a segmented trunk the change runs the other way: the split-mass spine takes the lead away from the arm at three segments, so	An endurance figure per architecture instead of per machine, and a cooling budget sized where the maximum is. The temperature itself needs per-actuator thermal data this review did not locate for the machine these numbers are calibrated against

	times-speed peaks highest even though the trunk carries more mass	the move is to stop shrinking modules	
Elastic pre-load saves 1 percent on a rigid trunk	Storage torque is 2E over the winding travel, so cost falls with travel and never with winding time	Long-travel, low-stiffness series elements wound through a segmented trunk's full range	A genuine stretch-shorten cycle, measured: at a 75 percent share the trunk actuator's peak draw falls 6.4-fold, 184 to 29 watts, with a mean of 8 watts relocated into the slow wind-up. The removal is at the actuator the element crosses, the component a drive is sized by; the whole-body peak stays with the arm
Momentum-scored strikes leave the torso still, at 8 percent of the copper bill	Effective mass and fist speed trade one for one in the score, and the arm is the cheaper channel	Multiply the score by a weight-behind-it gate, anneal the gate as a curriculum, and hold exploration open past the anneal; the additive and penalty spellings both converge to a shove, measured	Measured: the gated learner delivers 13.1 kg·m/s at 2.39 kg in four of five seeds, past this table's own target; momentum-only converges to a 7.5 m/s whip with 0.67 kg behind it
Compute is 1.5 to 2.5 percent of a strike, its arithmetic a millionth of that	Platform watts times seconds is the whole compute bill; the multiply-accumulates are noise within it	Duty-cycle the platform between control-relevant windows instead of shrinking the network	The efficiency budget moves to the terms that exist: trajectory shape, actuation topology, winding losses
The absolute scale rests on a calibrated reach, 71.5	This review did not locate a strike-energy measurement on humanoid	Pack-current metering through an instrumented strike against a force plate, on the	Joules per punch becomes a reported figure rather than a modelled one, and this table's first

J, rather than a measured strike	hardware; the model's scale is calibrated on a reach	platform already in the ring	row becomes testable on hardware

A humanoid trunk that pays for itself in a strike was **out of reach** while the only actuation on offer was one rotary module per joint, and this report prices exactly why. It is **demonstrably approachable today**: the multiarticular rows already sit within 5 percent of the single-joint baseline while carrying seventeen vertebrae of travel and compliance, and the tendon-driven humanoid literature has built such trunks for a decade. It becomes **ordinary** when trunk actuation is specified by actuator count and winding travel rather than joint count, which is a design-rule change, the most tractable kind of change there is.

> **Run it rather than read it.** The five benches and the seven run receipts every figure above is taken from are released under CC0 at huggingface.co/datasets/physicalai-bmi/punch-energetics. Single-file Rust, no dependencies: `rustc -O bench/punch_energetics.rs -o /tmp/pe && /tmp/pe` reproduces the 270.7 J baseline and the eight gates that precede it. Two cautions carried in the bundle's own README, because they decide what a comparison means: a gap under about five percent between two architectures is search scatter rather than a difference, and feasibility evidence pools one-sidedly, so a found strike stands while an absence means only that the draws attempted found none.

References

1. Filimonov, V.I., Koptsev, K.N., Husyanov, Z.M. and Nazarov, S.S. *Boxing: Means of increasing strength of the punch*. National Strength and Conditioning Association Journal 7(6), 65 (1985). Master-class decomposition: 38.5 percent leg drive, 37.3 percent trunk rotation, 24.1 percent arm extension. REPORTED · VIA THE SPORT-SCIENCE LITERATURE
2. Khatib, O. *A Unified Approach for Motion and Force Control of Robot Manipulators: The Operational Space Formulation*. IEEE Journal on Robotics and Automation 3(1), 43 (1987). The effective-mass definition used throughout. READ · FRAMING AND DEFINITION
3. Deniz et al. *Physics-based electrical power model of the Unitree G1 7-DOF arm*. arXiv:2606.15915 (2026). Copper, viscous, Coulomb and baseline-torque structure; R^2 0.965 on unseen speeds. The loss form and per-joint dominance pattern this bench reproduces. MEASURED · THIRD PARTY

4. Energy-aware reinforcement learning on the same arm, arXiv:2606.15918 (2026). Mean 71.5 ± 48.3 J per reach, measured on the physical G1. The calibration anchor for absolute scale. MEASURED · THIRD PARTY
5. Li, Yang, Chen, Chen, Lin and Li. *From Inference Efficiency to Embodied Efficiency*. arXiv:2603.19131 (2026). Methods that reduce computation under conventional metrics often increase end-to-end execution cost on robotic platforms. READ · ABSTRACT, VIA THE INSTITUTE RECORD
6. Unitree Robotics, G1 published specification: about 35 kg, 23 degrees of freedom in the base configuration, waist degrees of freedom by variant. VENDOR SPECIFICATION
7. CES 2026 humanoid boxing demonstrations, press record: Interesting Engineering, 6 January 2026, and the eWeek and Rest of World coverage of the 2025 Hangzhou tournament and the San Francisco exhibition bouts. REPORTED · PRESS
8. Seok, Wang, Otten and Kim. *Actuator design for high force proprioceptive control in fast legged locomotion*. IROS 2012. Motor-constant scaling for geometrically similar electric machines, the exponent Section 6 sweeps from 1.00 to 2.00. READ · SCALING ARGUMENT
9. Neto, O.P., Magini, M. and Saba, M.M.F. *The Role of Effective Mass and Hand Speed in the Performance of Kung Fu Athletes Compared with Nonpractitioners*. Journal of Applied Biomechanics 23(2), 139 (2007). Trained strikers present roughly twice the effective mass of novices at similar hand speeds; the basis for the 7 percent floor. REPORTED · ABSTRACT
10. Institute for Physical AI @ JBI, Charlot Lab. `punch_energetics.rs`, the joules-per-strike bench behind every table above: exact planar recursive Newton-Euler dynamics, the identified loss form of reference 3, a multi-start constrained search granted to every architecture, and the eight gates of Section 10 printed at the head of the run output alongside the full sweep. MODELLED · OURS, EIGHT GATES PRINTED
11. Institute for Physical AI @ JBI, Charlot Lab. `total_task_energy.rs`, the sibling bench: the same sum per completed reach, with the compression ablation and the calibration against reference 4. MODELLED · OURS
12. Institute for Physical AI @ JBI, *Energy Observability*, Technical Report TR-2026-36. The measurement doctrine this bench's gates apply. OURS · PRIOR REPORT
13. Institute for Physical AI @ JBI, *The Physical AI Hardware Lottery*, Technical Report TR-2026-40. States the whole-robot bound on any compute improvement as the reciprocal of the non-compute share, 1.03 to 2.44 times on the bodies measured so far; Section 9 here is the strike-side instance. OURS · PRIOR REPORT
14. Institute for Physical AI @ JBI, Charlot Lab. `punch_reward_learning.rs`: cross-entropy policy search over the same strike and loss model, thirteen reward conditions by five seeds, a parity gate that re-derives the 270.7 J baseline before any learning is trusted, and every prediction registered in the file header before the run it scores. Six of thirteen predictions fell; the falsifications carry the finding. Released CC0 in the punch-energetics bundle. MODELLED · OURS, PREDICTIONS REGISTERED
15. Institute for Physical AI @ JBI, Charlot Lab. `punch_demand_frontier.rs` and `punch_verify_c.rs`: the demand sweep with family-transfer feasibility and escalation rounds,

and the equal-budget replication behind Section 8's frontier and elastic figures, six independent draws per architecture per demand with every draw printed. Released CC0 in the punch-energetics bundle. MODELLED · OURS, REPLICATED

16. Lin, W., Qian, L., Luo, X. and Liang, C. *Temperature distribution prediction of the quadruped robot based on the lumped-parameter thermal networks*. Robot 47(2), 188–199 (2025), in Chinese. The in-situ identification Section 7 names: a whole-body thermal network least-squares fitted to measured motor temperatures, MSE below 2.5 $°C^2$. LOCATED · NOT READ, NO FIGURE TAKEN
17. Wan et al. *Learning to Balance Motor Thermal Safety and Quadrupedal Locomotion Performance with Residual Policy*. arXiv:2605.27046 (2026). Drives that identified thermal model into a locomotion policy; the source of the coupled-node and velocity-dependent-convection features Section 7 contrasts this report against. READ · ABSTRACT AND MODEL SECTION