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

SENSING · ENERGY · MEASUREMENT

Technical Report TR–2026–42
Research / Position · Preprint v1.0
20 August 2026

SENSING · ENERGY · MEASUREMENT

What a Machine Spends Looking

Task energy on a body is reported as compute against everything else. Sensing is measurable on the same meter, is not separated, and how large it has to be is set by the body before any perception exists.

1. What is in scope, and what does this continue?

One term of a sum that three companion reports have already established. TR-2026-40 §10 names the denominators a figure can be stated in, and shows that on a body the second of them decomposes: compute is one term, and improving it by a factor leaves the rest untouched, so the limit of that improvement is the reciprocal of the non-compute share. Measured, that share moves a long way: five percent of a reaching task at one watt of compute, fifty-nine percent at thirty watts, and on a mobile robot under autonomous navigation a graphics processor drawing 37.3 percent against 16.6 percent for the motors.

TR-2026-36 establishes where a meter has to sit before any of that is attributable, and grades *ISO/TS 25213*, the test method for the energy consumption of 6-axis articulated industrial robots, at stage 60.00 and effective 26 June 2026 CATALOGUE RECORD AND SCOPE VERIFIED; CLAUSE TEXT NOT READ. TR-2026-35 reads the transfer literature across twenty-three regions in eighteen languages, and TR-2026-43 measures the same body property this report sweeps, against demonstration volume rather than against a sensor set. The viability denominator this Institute reports against, joules per viability-second held under disturbance, is already demonstrated across six bodies including a physical SO-101 arm at 140 randomised seeds.

This report adds one term and one body property. The established split is compute against non-compute. Sensing sits inside the second half, undifferentiated, and it is separable on the same instrument. And the size of the sensing bill turns out to be set by a mechanical property that is fixed long before a perception stack exists.

> **Run it.** The three instruments behind this report are on the site and need nothing installed: Strip the senses sweeps pivot stiffness against a sensor set and is §3; The focus band recomputes an allocation's energy exactly under a different gate cost, slow-path cost and substrate idle floor; and The refusing receipt applies §7's conditions to a claim and names what would have to change in the measurement to fill each refused field.

2. What does a machine spend looking?

A term nobody separates, on an instrument that can already separate it. A sensor read costs energy in three places: the transducer's own draw, the conversion and transport of the sample, and the computation that consumes it. The first two are properties of the channel and are constant per read, which makes them straightforwardly meterable on a dedicated rail alongside compute. What makes the term worth separating is not its size but its *controllability*: unlike actuation, which a disturbed body must spend, and unlike a policy's compute, which is fixed by the architecture, the sensing bill is set by a schedule the designer chooses.

On the bench used here, four channels carry per-read costs of 0.0090, 0.0060, 0.0030 and 0.0025 joules for cart position, cart velocity, pole angle and pole rate. Priced against the region each buys, the attitude pair costs 2.20 J and buys its region at about 19.3 J per point; the position pair costs 6.00 J and buys it at about 164. **The protective channels are cheaper in absolute terms and by about 8.5x per unit of viability bought**, which is what makes an always-on protective layer affordable while continuous perception is not.

3. What sets the size of that bill?

The body, and it is decided in mechanical design. Sweeping the pivot stiffness of a cart-pole with every sensor switched off and the controller byte-identical in every row, the share of the episode held inside the barrier moves from **0.000** at zero stiffness to **0.360** at 0.50, **0.794** at 1.00 and **0.834** at 2.00 N·m/rad. Nothing about the policy changed. What changed is whether the upright is a stable equilibrium of the mechanism.

The threshold is a property of the body's own equations rather than of the textbook form. The gravitational stiffness of a pendulum about a *fixed* pivot is m g L = 0.4905 N·m/rad for this body, but the pole sits on a free cart whose recoil takes angular acceleration with it. Linearising the implemented step at the upright gives $m g L \cdot (4/3) / (4/3 - m_p/(m_c+m_p))$ = **0.5264**, 7.3 percent higher, confirmed by bisection on the discrete step at 0.5264. A design that labels stiffness against the fixed-pivot figure calls an unstable body stable across part of its range.

On a second body the same law holds, and the plateau sits higher. SENSE/2 runs the same sweep on a two-link standing body, ankle and hip, where the threshold is the largest eigenvalue of the coupled gravitational stiffness matrix at 51.25 N·m/rad, 1.13x the largest single-joint figure. The no-sensor region rises to **0.682** at that threshold, **0.984** by 1.5 times it and **1.000** by twice it, flat out to 2.7 times. An earlier version of this section reported that same curve peaking at the threshold and collapsing to zero past it, and concluded that the cart-pole's plateau was peculiar to its morphology. That collapse was the integrator rather than the body, and both halves are withdrawn in §7. A third body, which makes and breaks contact, separates this claim from its obvious confound in §5: the mechanical dial moves the sensing bill by a factor of four and the controller gains do not move it at all.

The complementary half is what sensing then buys. At a stiffness of 1.0 the blind loop holds region 0.794 while its task error sits at 0.936, indistinguishable from not attempting the task; the full four-channel suite takes task error to 0.487 and region to 1.000. **Sensing bought the task, not the survival.**

4. Why does more data not fix it either?

Because a dataset is a recording of sensor streams, so the two bills have the same floor. TR-2026-43 ran the same dial against a different resource on the same plant. Where §3 above sweeps pivot stiffness against a *sensor set*, that report sweeps it against a *demonstration count*, cloning a policy from camera observations the way a video pipeline sees the world. At maximum data, stiffness buys **+0.746** of certified operating region. At zero stiffness, **256 times more data buys −0.043**, flat within two standard errors. Two demonstrations on a passively stable body certify the full region; five hundred and twelve on a free pivot reach **0.254**.

Set beside §3 the two results say one thing. Compliance moves the region a blind controller holds from 0.000 to 0.834, and it moves the region a data-trained policy holds by +0.746, and in both cases the resource being scaled buys almost nothing without it. **That is not two findings that agree. It is one finding reached from two directions**, and the reason they cannot come apart is that a dataset is a recording of sensor streams. Torque exchanged directly between a compliant body and its load never crosses a transducer, so it is absent from the sensor stream in real time and absent from every recording of that stream, permanently.

The consequence is a shared floor, which neither report states alone. The sensing bill and the data bill are not two budgets to be traded against each other. They are two ways of paying for the same missing information, and both are bounded below by the same mechanical property. A programme that responds to a thin sensor budget by collecting more demonstrations is buying the second copy of what it already could not see. The dial that moves both sits in the mechanism, and it is set before either budget is written.

5. Does the result survive a change of body?

Three bodies now, and the third one changes what the question is worth asking about. A result from one body is a data point, so §3's sweep was repeated on a two-link standing body: two coupled rotations rather than a rotation and a translation, inertial coupling between the links, and no free base to absorb momentum. It has since been repeated on the body this report's own forcing function asked for, a hopping monoped, where a compliant leg carries the weight until the leg force reaches zero and at that instant the foot leaves the ground. The controller, the estimator, the disturbance, the seed schedule and the four sensor costs are held identical across all three.

Claim	Cart-pole	Two-link stand	Hopping monoped
The threshold is a coupled property, not a textbook figure	0.5264 against the fixed-pivot 0.4905, 1.07x	51.25 against the largest single joint 45.47, **1.13x**	64.7 N·m/rad, the gravitational stiffness about the foot; the sweep runs to 64x it
Compliance moves the no-sensor region	0.000 to **0.834,** monotone, flat from 2.0 out to 32	0.000 to **1.000,** monotone, flat from 2x the threshold	0.719 to **0.773,** monotone, and **1.000 is out of reach at every stiffness**
What sensing buys	the task: error 0.936 to 0.487, region already held	**the survival:** region +0.323, task error only +0.014	**almost nothing:** perfect state is worth 0.037 of region
Can the mechanism alone be tuned to sufficiency?	Yes, once the base has authority: 16 kg gives 1.000	Yes: 1.000 at twice the threshold	**No.** 64x critical stiffness with a perfect state estimate reaches 0.787

Across the two smooth bodies, two of the three original claims carry and one inverts. On the cart-pole the blind body survives and cannot do the job; on the two-link stand at its threshold stiffness the blind body is *marginal* at 0.677, and the sensors are what hold it inside the barrier. **Which of the two vectors a sense serves is a property of the body, not of the sense.**

The two smooth plateaus differ, and the difference is a second barrier rather than a disagreement. The cart-pole levels off at 0.834 and the stand at 1.000. Decomposing which barrier the blind cart-pole actually hits accounts for the gap. At a stiffness of 0.5 every failure is an attitude failure. At 1.0 it is 70 percent cart speed. From 2.0 out to 32, **not one failure is an attitude failure**: all of them are the cart running past its velocity bound. The compliant pivot has removed the falling problem completely and left a base-excursion problem in its place. Giving the base authority closes the gap, with the policy still unchanged: at 4 kg the plateau is 0.944 and at **16 kg it is 1.000**, matching the two-link stand, whose base is fixed to the ground and has no velocity bound to violate. **Stiffening a joint converts an attitude failure into a base failure, and whether that is progress depends on whether the base has somewhere to put the momentum.**

The sharper number on the stand is in the sensor sets. The **ankle pair alone**, at 2.20 J, reaches region 1.000 and task error 0.069: indistinguishable from the full four-channel suite at 8.20 J. The hip channels buy nothing measurable on that body. That is the same shape as §2's protective-against-generative pricing, at **3.7x** the saving, and it is only visible because the channels were priced separately.

The contact body confirms §3's central claim more sharply than either smooth body, and it does so by nearly running out of the quantity §3 measures. An oracle arm was added for it: the controller reading the true state, with no estimator at all, which is the ceiling any sensor suite could ever reach. At the stiffness the sensor sweep runs at, perfect, free, noiseless omniscience is worth **0.037** of certified region over knowing nothing. That is the entire prize every channel is competing for on this body.

That figure is a property of the body and not of the controller, which is the point. Across a fourfold change in position gain and a 2.4-fold change in rate gain, at two disturbance levels, six cells of 512 paired seeds each, it moves between 0.035 and 0.042: inside its own error bars. Turn the *mechanical* dial instead and it moves by a factor of nearly four, falling monotonically from **0.054** at zero ankle stiffness to **0.014** at 32 times the critical value. §3 asserts on one body that the mechanism sets the sensing bill. Here the two dials were turned independently over comparable ranges on a third body with different dynamics, and only one of them did anything.

What the contact body cannot be made to do is reach sufficiency. On both smooth bodies a stiff enough mechanism holds the barrier with the policy contributing nothing: the stand reaches 1.000, and the cart-pole reaches it too once its base has somewhere to put momentum. This body does not, at any stiffness, on either arm. The blind region plateaus at 0.773 and the *oracle* plateaus at **0.787**, and 64 times the critical ankle stiffness lifts neither. The reason is in the mechanism rather than inferred from the curve. The ankle spring acts through the foot; the foot is off the ground for 39 to 41 percent of every second; and in flight, angular momentum about the centre of mass is conserved, so no torque the machine commands changes its attitude. **The compliance is not weakened during flight, it is absent, and a mechanism that is missing for part of every cycle cannot be tuned into sufficiency.** Neither can a sensor. Roughly a fifth of the region is unreachable to a controller that knows everything.

Two channels chosen because this body has a discrete mode buy nothing, and the reason is worth more than the result. A single binary contact bit is the cheapest channel in the ledger at 0.0002 J a read, and it was the registered favourite: an estimator that is wrong about whether the foot is down is wrong about a category rather than about a number. It buys between −0.0005 and −0.0034 of region across all six cells, never positive, and in two of them resolvably below zero. Leg length is the same. What captures the headroom is the pitch pair at 0.041 to 0.048, and the **full four-channel suite captures less than the pitch pair alone in all six cells, at 1.8 times the joules**. That is §2's synergy result on a third body: per-sensor value is not additive, so a ranking that buys channels one at a time can buy the wrong set.

The mechanism is that a partial correction to a hybrid state is worse than no correction. Reading pitch alone leaves the estimator in the wrong contact mode on 4.25 to 8.96 percent of steps against a blind estimator's 1.40 to 2.38, a gap of 2.57 to 6.91 points resolved in all six cells. Correcting the angle without its rate desynchronises the estimator's prediction of when the foot lands. The contact bit does the same thing from the other side: forcing the estimated mode to the truth while the estimated leg state still disagrees makes the next switch worse, and it raises mode error in five of six cells, resolvably in three. Adding the rate collapses the error to 0.35 to 1.13 percent, resolved in all six. **On a body with a discrete mode, sensing one component of the state without the components that determine the switch is a negative purchase.**

One reading in that sweep is not a claim. The pitch pair reads at 111 to 121 percent of the headroom in all six cells, which is to say it appears to do slightly better than the perfect state estimate it is approximating. With a fixed sub-optimal controller nothing forbids that, and a 128-seed version of this bench read 132 percent. The paired difference at 512 seeds is 0.006 ± 0.010 of region, so it is not resolved here either. What the larger sample adds is a bound rather than an answer: whatever the effect is, it is smaller than 0.016 of certified region, which is smaller than the headroom it sits inside. **The quantitative form of the ceiling is also not claimed.** Varying hop energy to change the airborne fraction moved the gap non-monotonically at 64 seeds, so this report claims that the ceiling exists and that the mechanism sets its height, and does not claim that the airborne fraction predicts it.

Read together the three bodies say something the first two could not. §3 measures how much a body makes its machine spend looking. The contact body's answer is 0.037, and the honest way to state it is not that sensors are cheap here but that **the ceiling which makes them cheap is the same ceiling that caps the machine**. Information about a state the machine cannot act on has no price, and for two fifths of every second this machine cannot act. **The dial that lifts that ceiling is not in the sensor budget and not in the controller gains. It is an internal momentum device**, a reaction wheel or a coordinated limb swing, which is the only thing that gives a body authority over its own attitude while it is off the ground. SpaceBok carried a reaction wheel for exactly this reason, and later work replaces the wheel with leg motion at the cost of a harder controller. That is mass and energy spent to make a sensor worth buying, and §11 carries it.

§11's forcing function named a dial for that ceiling, and the dial has since been turned. That row was an assertion about a mechanism, so it was tested rather than left standing. A reaction wheel was added to the contact body and *only its torque limit* was raised, with the compliance, the disturbance, the controller, the seed schedule and the airborne fraction all held fixed. The device is specified by the two numbers that bind a real one: the torque it can apply, and the angular momentum it can store before it saturates and can no longer push the same way.

Wheel torque limit, N·m	Blind region	Oracle region	Headroom	Airborne	Wheel copper, J

0, the body §5 reports	0.736	0.775	0.039	36.3%	0.0
10	0.736	0.813	0.078	36.3%	67
20	0.736	0.852	0.117	36.3%	271
40	0.736	0.933	0.197	36.3%	973
80	0.736	**0.985**	0.249	36.3%	2041
160	**0.736**	**0.988**	**0.252**	36.3%	2344

The ceiling was absent authority and not the contact. The oracle rises from 0.775 to 0.988 while the foot is off the ground for the same 36.3 percent of every second throughout, and the result holds in all three cells run: 0.775 to 0.988, 0.777 to 0.988 at twice the position gain, and 0.846 to **1.000** at the lower disturbance, where perfect state plus authority solves the task outright exactly as it did on both smooth bodies. A body that makes and breaks contact behaves like a smooth one again as soon as something can act while it is airborne. Running the wheel in stance as well changes almost nothing, 0.985 against 0.988, so the flight-only restriction isolates the variable rather than producing the effect.

And every bit of that accrues to the arm that can see. The blind column does not move: 0.736 at a torque limit of zero and 0.736 at 160, across a range in which the oracle gains 0.213 of region. The headroom, which is everything any sensor could ever buy on this body, rises by a factor of **6.4**, from 0.039 to 0.252, and by 5.0 and 6.1 in the other two cells. **Authority and information are complements rather than substitutes, and neither is worth much without the other.** A machine that cannot act gains nothing from a sensor, and a machine that cannot see gains nothing from an actuator.

The mechanism is measurable and is not a matter of interpretation. The blind arm's *delivered* impulse is flat at 0.32 N·m·s whether its limit is 5 or 160, while the oracle's rises from 2.00 to 20.84. The blind controller never asks for the torque the device could give, and the reason is in the two error columns: its true attitude error is 0.0830 rad while the error in the estimate it is regulating is 0.0254, a factor of 3.3. **A blind estimator's error is invisible to itself, so it believes it is close and**

asks for almost nothing. A torque limit amplifies a command; it does not create one.

What the dial costs, with the grade its constant carries. The wheel draws 2,344 J of copper across an episode at the top of the sweep, against a sensor suite that costs 0.10 to 4.85 J: roughly five hundred times, and twenty thousand times the contact bit alone. That figure rests on a motor constant which is *calibrated rather than measured*, exactly as the four channel costs are, and no number derived from it should be quoted without that grade. What does not depend on the calibration is the ordering and the shape: **the device that makes sensing worth buying costs orders of magnitude more than the sensing it makes worth buying**, and it is bought in mass and joules before a perception stack exists. That is §3's claim arriving from the actuation side.

One registered prediction was refuted, and the harness that would have hidden it is worth naming. The expectation was that the contact bit would finally earn its place, because a wheel that fires only in flight must decide *when*, and it decides on the estimated contact. It does not: the bit buys −0.0005 to −0.0027 of region at every torque limit in every cell, never positive. The first version of this experiment could not have asked the question at all. The flight gate sat inside the integrator, which sees the true contact flag of the body it is integrating, so the actuator was reading ground truth for free. That is the third time this family of benches has leaked the true state into something that drives the body, and moving the gate to the estimator's side is what made the prediction testable. It was still refuted.

6. Which hypotheses does this test?

Four, stated before the bench existed, and one survived.

#	Hypothesis	Resolution
H1	The zero-sensor certified region is nonzero	**Survives, and sharpens.** It is a property of the body: 0.000 at zero pivot stiffness, 0.834 at 2.00, policy unchanged.
H2	Sense value is strongly sublinear; the first sense	**Falls.** Not sublinear and not monotone: the second channel buys more region than the first (+0.053 then +0.061), so value here is synergistic rather

	buys most of the region	than additive and a per-sensor ranking can mis-order the set to buy. Whether attitude alone helps or hurts tracking is *not resolved*: +0.006 ± 0.030 at 512 seeds, and no direction is claimed.
H3	Protective sensing costs more per unit of region and less in total	**Falls, in the useful direction.** Cheaper on both counts, by about 8.5x per region point.
H4	Sampling triggered by the agent's own intervention beats a clock	**Falls outright.** Scored against a no-sensor control the trigger is inert, holding to within 0.002 of what drawing zero samples holds at every disturbance level.

The contact body registered two more before it ran, and the one that mattered inverted. Its second hypothesis was that an intermittent preflex cannot be tuned to sufficiency, because a mechanism that is absent for part of every cycle cannot be made to hold a barrier on its own. It survives: 0.787 at 64 times the critical ankle stiffness with a perfect state estimate, against 1.000 on both smooth bodies. Its third was that a single binary contact bit would be the best buy per joule on this body, because a mode error is categorical where a drift is not. **It is refuted by the most tightly resolved number in the sweep**: the contact bit buys −0.003 to −0.000 of region across six cells and is never positive. The inversion is the better result, because the question of which channel is the best buy presumes there is something to buy, and §5's headroom says there is almost nothing.

The wheel experiment registered four more and one fell. That the ceiling would lift with authority, that the headroom would grow with it, and that the wheel would cost far more than the sensing all survive, in three cells each. That the contact bit would finally be worth buying does not: it buys −0.0005 to −0.0027 of region at every torque limit, in every cell, and is never positive. Across §5 the cheapest channel in the ledger has now been offered three separate reasons to matter on the one body whose state is discrete, and has bought nothing every time.

7. A result of ours that was wrong, and how it was caught

Four. Two were errors of method, the third was an error of numerics, and the fourth was a citation that had gone stale.

The first version of this bench reported the attitude-alone tracking gap as **1.078 against 0.973**. Per-episode spread of task error is about 0.25, so sixteen seeds resolve nothing finer than 0.06 while the effect is 0.03, and roughly a third of sixteen-seed schedules invert its sign. The direction survives at 512 seeds; the magnitude was about three times too large and is withdrawn. **What moves it:** the standard error is now printed beside every mean, so the size of an effect and the resolution of the sample are visible in the same glance.

The second reported a conditional law: efference-triggered sampling loses under disturbance and wins in a quiet world on a twentieth of the clock's samples. It does. So does drawing no samples at all, on a body that needs none. **The comparison lacked a no-sensor arm**, and the harness additionally granted the estimator a free exact state read on every barrier exit, which kept the trigger alive enough to resemble a policy. With the control restored and the read removed, the trigger is inert. **What moves it:** every sensing policy is now scored against the no-sensor arm, which turns "cheaper than the clock" into "worth more than nothing".

The third is the two-link body's stiffness sweep in §5, and it was the integrator rather than the harness. Forward Euler on a spring-damper is stable only while the step obeys $DT \cdot \omega < 2\zeta$, and substituting $\zeta = C/(2\sqrt{(kp \cdot I)})$ and $\omega = \sqrt{(kp/I)}$ the inertia cancels and the whole condition becomes $kp \cdot DT < C$: a bound on *stiffness*, which is the quantity being swept, and which no inspection of the plant makes visible. That capped the two-link body at 45 N·m/rad and the sweep ran to 138. Above the bound the integrator injects energy on every step and the divergence reads as physics: the region was published as rising to **0.678** at the threshold and collapsing to **0.000** by twice it, and a conclusion about morphology was drawn from that shape. The guard in place at the time bounded the sweep by $DT \cdot \omega < 1.4$, which is the criterion for a *critically* damped mode; this family is lightly damped, so that guard reported a ceiling of 344 where the true one was 45. Substepping the physics while holding the controller at 50 Hz removes the collapse entirely: 0.241 becomes 0.984 at 1.5 times the threshold and 0.000 becomes 1.000 at twice it, converged against a 32-substep control. There is no optimal compliance and no disagreement between the morphologies. The cart-pole was 1.33 times over its own ceiling at the top of its sweep, where the error was small enough to hide, and the published instrument let a reader drag the slider there. **What moves it:** each bench now derives its ceiling from its own damping constant and refuses a sweep that reaches it, and a preflight gate re-integrates every plant at four times the substep count and fails if a published region moves. The gate's self-test restores the original substep count and requires the defect to reappear.

What survives the first two is the same sentence, on firmer ground: the efference copy tells a machine what to *subtract*, not when to *look*. In this model the agent predicts its own contribution exactly, so its own actions are never the dominant source of prediction error, and there is nothing for an effort-triggered sampler to track.

The fourth is §4, and none of it was our own measurement. Every figure §4 attributes to TR-2026-43 was quoted from a run made before that report's substepping fix: **+0.742** where the report publishes +0.746, **−0.027** where it publishes −0.043, and **0.258** where it publishes 0.254. The companion report was right throughout and had been corrected by the same integrator fix described above; this report went on quoting the superseded run, in the section immediately before the one that explains the fix. The description drifted with the numbers, calling a free pivot "a stiff one". **What moves it:** a figure this report attributes to another is now checked against that report's published text rather than against a run, because a cited number's source of truth is the cited document. The three replaced values are the gate's own control: planted back, it names all three.

8. What is this report careful not to claim?

Actuation energy here is physical, integrated as torque against velocity with a copper-loss term along the trajectory each policy actually produced. **Sensing and compute energy are parameterised**, calibrated so that a full always-on suite lands within an order of magnitude of actuation, which is the regime where the allocation question is interesting. Absolute joules are indicative and every result is a comparison at fixed calibration.

The body is a cart-pole with a lumped pivot spring-damper standing in for intrinsic actuator impedance. Whether the ratios here transfer to a legged or multi-contact body is not established by this report and should not be assumed in either direction. The region reported is a rescaled barrier-exit count over seeded rollouts, not a proof of set invariance; where this report says region it means that quantity, and the viability denominator this Institute reports against remains joules per viability-second.

9. What would sharpen this?

A body where the compute rail is separable on a shunt. Everything above is simulation with one physical term. The single measurement that would convert it is a small compliant body whose compute draws through its own instrumented rail, so that the three terms can be read apart rather than modelled apart. TR-2026-36 establishes the condition; this report supplies the reason to spend it on a sensing question.

Both terms on one meter. Authority has now been varied on its own and is in §5: a reaction wheel lifts the ceiling from 0.775 to 0.988 and multiplies the headroom by 6.4, while the blind arm does not move at all. What that leaves open is the exchange rate. The wheel's 2,344 J and the sensor suite's 4.85 J are computed from two different calibrated constants and never met a physical instrument, so the five-hundred-fold ratio between them is a modelling result rather than a measurement. The body TR-2026-36 specifies, with a compute rail and a sensor rail read separately, would settle it if its actuator rail were read too.

A fourth body where the airborne fraction is the independent variable. §5 shows that the ceiling exists and that ankle compliance moves the headroom under it by a factor of four. It does not show what sets the ceiling's height. The obvious candidate is the fraction of time the machine has no authority, and the obvious experiment, varying hop energy, moved the gap non-monotonically at 64 seeds and is not reported as a result. A body whose duty cycle can be set directly, rather than through an actuator that also changes the disturbance response, would settle it.

A per-read cost taken from datasheets rather than chosen. The four channel costs used here are calibrated, not measured. Real transducer and conversion figures would make the joules-per-region ranking a statement about hardware rather than about a parameterisation.

10. Conclusions

Sensing is the term of the task-energy sum that nobody separates, and it is the one term whose size the designer directly chooses. On the body measured here it is also the term whose necessary size is fixed by mechanism: a compliant pivot moves the region a blind controller holds from nothing to four fifths, and the sensing suite is left buying the task rather than the survival. The threshold that governs it is a property of the body's own equations and sits 7.3 percent above the textbook figure.

Read beside TR-2026-43, which sweeps the same body property against demonstration volume, the result generalises past sensing: compliance buys +0.746 of certified region there while 256 times more data buys −0.043. A dataset is a recording of sensor streams, so what never crosses a transducer is missing from both, and the two budgets share a floor set in the mechanism.

Two of the four hypotheses stated here were wrong in ways a control arm and a standard error would have caught. A third correction was not a hypothesis at all: an integrator that could not represent the stiffnesses being swept, which turned a monotone curve into a peak and a collapse and very nearly turned that into a claim about morphology. All three are withdrawn in §7 with the change that prevents a recurrence. The methods that produced those corrections cost less than the corrections did.

11. The forcing function

What is missing	Why	What builds it	What becomes possible
Sensing separated	The established split is compute against non-compute; sensing sits	A compliant body with the compute and sensor rails	A sensor suite chosen by joules per unit of

on a real meter	inside the second half undifferentiated	instrumented separately	viability rather than by capability, before the perception stack is written
The exchange rate between a joule of authority and a joule of sensing	§5 now measures both ends and they are not commensurable by anything published here: a wheel that lifts the ceiling from 0.775 to 0.988 draws about five hundred times the sensor suite it makes worth buying, and that ratio rests on a motor constant this report calibrated rather than measured	A body whose wheel and whose sensor rail are both instrumented, so the two are read on one meter rather than modelled against each other	A bill of materials in which mass, authority and sensor count are traded in a single currency, which is the form a design decision actually takes
Authority as a continuous design variable on a body that walks	The wheel result is measured on one hopping monoped whose duty cycle cannot be varied. Changing hop energy moves the disturbance response with it; gating the hop to every tenth touchdown moves nothing, because this body makes about three flight excursions in an episode and three events do not resolve a duty cycle	A body that leaves the ground *often*, in many short flights rather than a few long ones, so that a duty cycle exists to command at all	A statement of how much of a machine's sensing budget is decided by its gait rather than by its hardware
Sensor sets priced on bodies with a discrete mode	On the contact body the two channels chosen because the body switches, a contact bit and leg length, buy nothing, and reading pitch without its rate makes the switch estimate worse than reading nothing at all	The same sweep on multiple contacts, on swing-leg placement and on slip, where a mode error is a fall rather than a delay	A perception specification for walking machines that names which partial state reads are worse than no read, before the estimator is written
Per-read costs from	The channel costs here are calibrated rather	Transducer and conversion	A ranking that is a statement about

datasheets	than measured	figures for a real sensor set	available hardware
The shared floor stated as a design rule	The sensing bill and the data bill are treated as separate budgets to trade against each other, when both are bounded below by the same mechanical property	Sweeping compliance against a sensor set and a demonstration count on one body, which TR-2026-42 §3 and TR-2026-43 §5 now do separately	A programme that stops answering a thin sensor budget by collecting more demonstrations, because it can see that both buy the same missing information
A course that teaches the loop's energy	Energy is taught as a datacentre or trajectory-planning subject; embodiment as perception and policy	PAI-330, twelve labs, benches that can fall	Engineers who price a sense before specifying it, and who write a disclosure that refuses what their setup cannot support

References

1. Institute for Physical AI @ JBI, *Energy Observability in Embodied Systems*, Technical Report TR-2026-36. COMPANION REPORT, READ IN FULL
2. Institute for Physical AI @ JBI, *The Physical AI Hardware Lottery*, Technical Report TR-2026-40, §10, three denominators and the non-compute bound. COMPANION REPORT, READ IN FULL
3. Institute for Physical AI @ JBI, *Sim to Real: The Road to Physical Agency*, Technical Report TR-2026-35, twenty-three regions in eighteen languages. COMPANION REPORT, READ IN FULL
4. Institute for Physical AI @ JBI, *The Two Operating Systems: What a Body Carries That No Dataset Contains*, Technical Report TR-2026-43, §5, the substrate-against-data sweep on this plant. COMPANION REPORT, READ IN FULL
5. ISO/TS 25213, *Robotics: Test methods for measuring the energy consumption of robots: 6-Axis articulated industrial robots*, ISO/TC 299, ICS 25.040.30, stage 60.00, effective 26 June 2026. CATALOGUE RECORD AND SCOPE VERIFIED; CLAUSE TEXT NOT READ
6. E. von Holst and H. Mittelstaedt, "Das Reafferenzprinzip: Wechselwirkungen zwischen Zentralnervensystem und Peripherie," *Die Naturwissenschaften* 37, 464–476, 1950. READ IN TRANSLATION
7. W. Zhu et al., "Identifying Important Sensory Feedback for Learning Locomotion Skills," arXiv:2306.17101. READ IN FULL

8. D. Attwell and S. B. Laughlin, "An Energy Budget for Signaling in the Grey Matter of the Brain," *J. Cereb. Blood Flow Metab.* 21, 1133–1145, 2001. READ IN FULL
9. M. Papadakis, J. A. Olsen, I. Poulakakis and K. Alexis, "Modeling and In-flight Torso Attitude Stabilization of a Jumping Quadruped," arXiv:2409.14567, 22 September 2024. Torso orientation is tracked by generating corresponding leg motions rather than by a wheel. ABSTRACT READ; METHOD AND EXPERIMENTS NOT READ
10. Robotic Systems Lab, ETH Zurich, *SpaceBok*, 2018–2021, a 22 kg quadruped whose modularity "allowed hardware changes for several experiments, such as feet for granular soil or a reaction wheel for in-flight attitude control". PROJECT PAGE READ; HARDWARE REPORTS NOT READ
11. Institute for Physical AI @ JBI, Charlot Lab, the SENSE/1 sensor-ledger bench, the SENSE/2 two-link bench, the SENSE/3 contact bench and the FOCUS/1 allocation bench, 512 paired seeds per arm on the sensor sweeps and 16 on the allocation bench, dispersion reported on every mean, versioned at `datasets/loop-energy/bench/`. THIS REPORT'S INSTRUMENTS; EVERY FIGURE REPRODUCED FROM THEM