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

SECURITY · ENERGY · MEASUREMENT

Technical Report TR-2026-36
Research / Position · Preprint v1.1
13 August 2026

SECURITY · ENERGY · MEASUREMENT

Energy Observability in Embodied Systems

What the meter can prove about a machine, and where. Morphology decides whether computation shows up in the joules at all, and on the bodies Physical AI actually uses this review did not locate a method that settles it, which makes the instrument itself the opportunity.

The Charlot Lab · Institute for Physical AI @ John Bailey Institute
Evidence assembled 12 to 13 August 2026.

6 hypotheses resolved · 2 of our own claims refuted · 1 novelty claim withdrawn · 5 primary sources read in full

Software asserts; energy is spent. A power trace is the one account of a machine's activity that its own code cannot simply write, which makes energy a candidate substrate for a machine that must demonstrate its behaviour rather than claim it. This report asks how far that demonstration reaches on an embodied system, and answers from measurement rather than from architecture. Three findings carry it. First, the compute-to-actuation ratio runs opposite to the common assumption: on an autonomous mobile robot the motors draw 16.6 percent of system power while the GPU alone draws 37.3 percent [1], and the same chassis driven manually with autonomy disabled puts the motors at 64.0 percent, so the workload sets the ratio and autonomy is the case that matters. Second, a physics-structured power model of a humanoid arm predicts electrical power to 3.58 W RMSE on operating points absent from identification, against a 121.7 W full scale, but only 10.1 percent of that power is mechanical work: the dominant terms are copper loss and viscous friction, and they are fitted, not derived [2]. Third, and most consequential for any claim built on an energy balance, 104 of 1,017 trajectories in that same study had to be discarded because the robot's own balance controller leaked leg energy onto the arm's measured rail, a confound invisible to the electrical measurement and detected only with an inertial sensor [2]. No adversary is involved in that failure; it is native to a body whose parts push

on each other. Six hypotheses are stated and each resolved to a position, a binding constraint, and a threshold that would move it. Two claims made in this Institute's own working drafts were refuted by the evidence and are reported here at the same size as the findings, and one novelty claim was withdrawn on locating prior art. The binding constraint throughout is engineering implementation: commodity platforms expose a battery gauge rather than per-rail energy. The one published measurement method covers six-axis industrial arms, is voluntary, and excludes by name the morphology where the energy signal is most informative.

1. What is the difference between asserting and spending?

Software asserts; energy is spent. A power trace is the one account of a machine's activity that its own code cannot simply write, which makes energy a candidate substrate for a machine that has to demonstrate its behaviour rather than claim it. Every claim a machine makes about itself passes through the machine. A log entry is a statement its software chose to write; a signed receipt is a statement its software chose to sign. If the software is wrong, or has been made wrong, the statement is wrong in a way that is internally consistent and leaves no trace in the record that produced it. This is not a hypothetical failure mode so much as the reason attestation exists as a field.

Energy is different in one specific respect: **you cannot spend a joule by asserting it.** Dissipation is obligated rather than optional, which is why power analysis has been a working technique for decades in both directions, and why an energy trace is a candidate for evidence in a way a log is not. The Institute already prices decisions in joules [6] and already ships controllers that carry their own stability certificates [7]. The natural question is whether the same quantity can carry an integrity claim on a body.

This report does not answer that question with an architecture. It answers with three measurements, and the answers are narrower and stranger than the architectural argument predicts.

2. What is in scope, and how was the evidence graded?

The scope is deliberately small. This is a position report over a small, fully-read evidence base, not a survey. Three primary documents were retrieved and read in their entirety: a power-profiling study of an autonomous mobile robot [1], a parameter-identification study of a humanoid arm's electrical power [2], and the catalogue record and scope statement of the one published measurement method for robot energy [3,4]. One vendor specification is used for a device power envelope [5].

Every figure below carries a grade. **Verified** means the primary document was retrieved and the number read from it. **Reported** means a named source states it and this review did not confirm it independently. **Vendor** means a manufacturer's own published specification. **Modelled** means our arithmetic over graded inputs. Where a quantity was sought and not found, the text says that this review did not locate it and names the search; that is not a claim that the quantity is unpublished.

> **A note on what this report is not.** A companion working brief within the Institute's orbit treats the general architecture of energy-rooted attestation, its global standards landscape, and its hardware-root-of-trust position. That document belongs to a commercial programme and is not the subject of this one. Results owned there are cited as premises where needed and are not restated as findings here. This report carries only what the embodiment adds, which is the term the general argument does not have: a body.

3. Which hypotheses does this test?

Six hypotheses, each resolved in section 10 to a position on the trajectory, the constraint currently binding it, and a measured change that would move it.

#	Hypothesis	Status after this review
H1	What a measurement can attest is bounded by the plane it sits on	supported
H2	Actuation supplies a derived expectation, unlike general compute	refuted as stated; survives weakened
H3	An energy balance does not close on a body	supported, measured
H4	Custody assumptions fail when the machine can be taken	open

H5	A receipt computed by the host is not evidence	holds by construction
H6	The disclosure is the task, not the model	prior art located; a port, not a finding

Table 1. Two of the six did not survive contact with the evidence in the form they were first written. Both are reported below at the size they were claimed.

4. What is the compute-to-actuation ratio, measured?

The intuition is that a robot is a machine for moving, so movement is where the energy goes and the computer is a rounding error. It is worth stating plainly because this Institute's own first draft of this work was built on it, and it is wrong for the case that matters.

Liu, Shi and Shin profiled an autonomous mobile robot across four scenarios, collecting power per component from the sensors, the motors and the Jetson board, with the board's internal breakdown read from tegrastats [1]. Under manual control with the navigation stack and object detector disabled, the motors draw **64.0 percent** of system power. Enable navigation and detection and the motors fall to 32.0 percent while the GPU rises to 30.4 percent. Under autonomous control the motors draw **16.6 percent** and the GPU alone draws **37.3 percent**. The authors state the consequence directly: the cyber subsystem dissipates more power than the sensors and motors. Total system power ranges from 30 to 90 W, and navigation plus detection together account for 30 W. VERIFIED

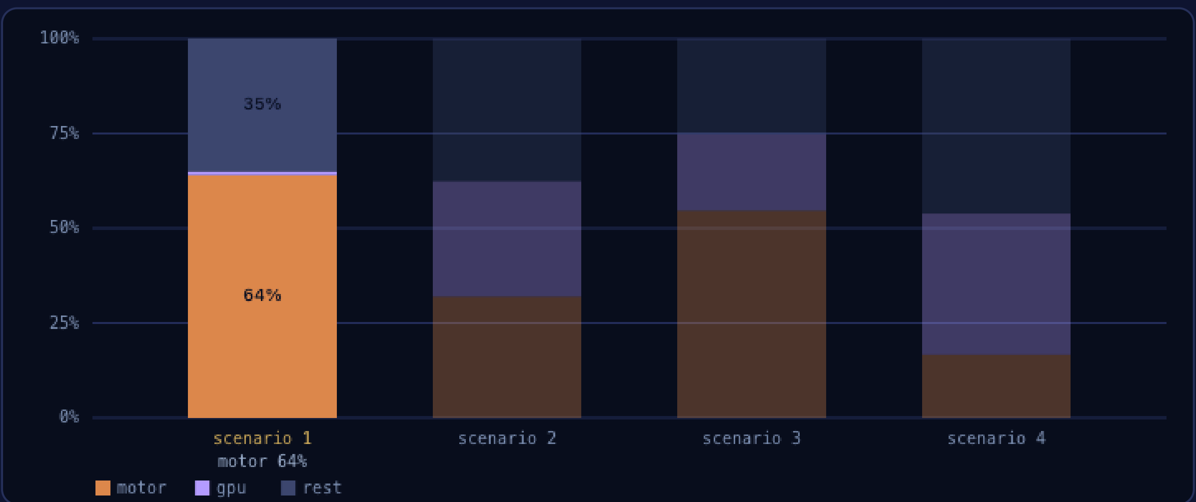

`1 · manual, no autonomy` `2 · manual + nav + detect` `3 · faster, + nav + detect`

`4 · autonomous`

`Scenario 1 · manual, autonomy off:` `motors 64.0%, GPU 1.0%. Nothing running but the wheels. This is the picture the masking intuition comes from.`

Figure 1. The same chassis, four workloads. Nothing about the hardware changes across these bars; only what the robot is asked to do. Predict before selecting scenario 4: the motor share there is lower than most readers expect by roughly a factor of four. **Only the motor and GPU shares are read from reference 1, figure 3** `VERIFIED`; the balance is drawn as one undifferentiated remainder `DERIVED` because the sensor, CPU and other components were not read individually and are not invented here.

The design consequence is immediate and it is the opposite of the masking argument. On a mobile platform under autonomy, computation is *the larger* term in the energy budget, so a whole-system meter is well placed to see changes in computational work rather than badly placed. The masking intuition is not a fact about robots; it is a fact about a robot being pushed around by hand.

> **A claim of ours, refuted.** An earlier internal draft of this work asserted that actuation dominates compute "by one to two orders of magnitude" on a mobile embodied system, and built a masking argument on it. Against reference 1 that is not merely unsupported, it runs the other way: the GPU-to-motor ratio under autonomy is about 2.2, and a widely-used robotics compute module is specified at 15 to 60 W [5] against the 135 W whole-robot rest draw of a humanoid [2]. The masking argument is withdrawn.

5. How well does physics predict an embodied system's power?

If a machine is to demonstrate what it did, there must be an expectation for it to match. For general-purpose computation that expectation has to be learned from a workload baseline, which is the structural weakness of any detector built on it: the estimator's model becomes the distribution being tested. Mechanical work invites a better answer, because a known mass moved along a known trajectory against known resistance has a computable cost. The question is how much of the electrical power that actually accounts for.

Deniz, Vega, Parsons and Auat Cheein identified a linear-in-parameters electrical power model for all seven joints of a Unitree G1 humanoid arm, using the robot's own onboard power sensor as the regression target [2]. Per joint, the model carries a mechanical term, a baseline-corrected copper-loss term, and Coulomb and viscous friction terms, augmented with pairwise joint-speed interaction terms: 49 parameters in total, 28 per-joint and 21 pairwise. It was fitted on 897 trajectories after outlier filtering and validated on 46 further trajectories executed at four speeds absent from identification.

Quantity	Value	Grade
Identification fit, 897 trajectories	R^2 0.933 · RMSE 1.07 W · MAE 0.86 W	VERIFIED
Hold-out at unseen speeds, 46 trajectories	R^2 0.965 · RMSE 3.58 W · MAE 2.33 W	VERIFIED
Hold-out error relative to full scale	2.9 % of 121.7 W	VERIFIED
Worst single-trajectory residual	17.2 W	VERIFIED
Mechanical work, share of predicted power	10.1 %	VERIFIED
Copper loss · viscous friction	42.4 % · 38.2 %	VERIFIED
Whole-robot rest power	135 W incl. legs and computing boards	VERIFIED
Arm rail baseline	120 W, excludes most leg and computer loads	VERIFIED
Power sampling rate	≈1 Hz, against 100 Hz joint states	VERIFIED

Table 2. The identified model and its error, from reference 2. The hold-out row is the one that matters: it is the error at operating points the model never saw.

Two readings follow, and they pull in opposite directions. The optimistic one is that the model generalises: an R^2 of 0.965 at speeds absent from identification is better than the 0.933 achieved on the identification set itself, which is the signature of a model whose structure is right rather than one that has memorised its training distribution. That is exactly the property a black-box baseline cannot offer, and it is the defensible form of the claim.

> **A second claim of ours, substantially wrong.** The same internal draft described actuation energy as "derived from first principles". Only **10.1 percent** of the arm's electrical power is mechanical work. Copper loss and viscous friction together account for 80.6 percent, and their coefficients were fitted from 897 trajectories, not derived from a datasheet. the description is *physics-structured with identified parameters*. The structure is physics; the magnitudes are measurement. Any argument that leaned on "derived" has to be re-run with "identified", and the conclusion weakens accordingly.

The practical number for anything built on this is the residual. A 3.58 W RMSE sets a floor: work costing less than roughly three standard deviations of that residual is inside the model's error, and the correct word for it is invisible rather than absent. On the worst trajectory the residual reaches 17.2 W. The authors also note that per-sample fitting is numerically degenerate on this data, with an R^2 below 0.04 and a condition number around 7.8×10^{11}, so trajectory-level aggregation is a requirement rather than a convenience. VERIFIED

6. Where does the balance fail to close?

The most useful result in this report is one nobody set out to produce, and it is a measurement of the confound rather than of the signal.

Of 1,017 trajectories collected, **104 were rejected** for what the authors term balance-compensation contamination [2]. During fast or multi-joint arm manoeuvres, the robot's locomotion stabiliser adjusts leg posture to compensate the arm's dynamic reaction. That leg work is real energy, and on this platform some of it lands on the same rail through which the arm is measured. The arm's power reading therefore includes work the arm did not do.

Nothing in the electrical signal announces this. The reading is plausible, the residual is unremarkable, and a model fitted through the contamination quietly fits the wrong thing. What exposed it was a **second and physically different sensing modality**: trajectories whose body angular-velocity standard deviation exceeded 2.5 times the median were rejected on that basis. Removing them moved the residual kurtosis from 77 to 2.97, which is to say from a heavy-tailed distribution to a nearly Gaussian one. VERIFIED

> **Why this belongs in a Physical AI report rather than a security one.** There is no adversary in this finding. A body whose parts exert forces on one another moves energy across subsystem boundaries as a matter of mechanics, and a rail-level meter cannot distinguish work that crossed the boundary from work done locally. At a measured 10.2 percent of trials on a current commercial humanoid, this is not an edge case. Any energy-balance argument applied to a body inherits it, and inherits it before anyone attacks anything.

The general rule the finding supports is one the attestation literature reaches by a different route: never verify an energy balance with a single measurement modality. Two physically distinct mechanisms, for instance an electrical measurement and an inertial or thermal one, are much harder to fool deliberately. This study demonstrates that they are also much harder to fool *accidentally*, which is the more common case and the one that silently corrupts a dataset.

7. How do you verify the instrument before believing it?

A measurement plane can be entirely dead while producing confident, plausible output, and that failure never announces itself. In this Institute's own instrumentation record, a laptop battery gauge read **67.2 W** at idle and **67.2 W** with ten processor cores pinned at full load. Not approximately: identically. VERIFIED, OUR OWN MEASUREMENT

Zero variance under a load you control is the signature. It matters more than a noisy instrument does, because a monitor built on such a plane reports healthy for every input indefinitely, and is never wrong in a manner anyone notices. A noisy instrument produces suspicious results; a dead one produces reassuring ones.

Separately, and on a different plane, run-to-run variation in this lab's GPU measurements proved **bimodal and whole-run**, driven by clock state rather than per-sample noise. A per-sample threshold detector is mis-specified against that structure; ratios taken against a simultaneous control cancel it, absolute thresholds do not. VERIFIED, OUR OWN MEASUREMENT

The reporting rule this Institute therefore proposes for any energy-based claim, and applies to itself in this document, has three required fields: **the measurement plane**, **the instrument's noise floor against a known load**, and **the null control**. A ratio reported without a null control is not a result. Where within-item spread equals or exceeds between-item signal, the action is to refuse to rank rather than to rank quietly.

8. Who is entitled to define a method?

Until recently there was no agreed procedure for stating how much electricity a robot uses, which is why vendor figures were never comparable. **ISO/TS 25213**, *Robotics: Test methods for measuring the energy consumption of robots: 6-Axis articulated industrial robots*, addresses that for one morphology. It reached stage 60.00, standard under publication, under ISO/TC 299, with an effective date of 26 June 2026 [3], developed with the Swedish Institute for Standards and participants from eleven countries [4]. Its scope specifies the measurement methods, the conditions for the measurements, and how results are presented. It **excludes service robots, medical robots, SCARA, autonomous mobile robots and delta robots**. Humanoids are not named. VERIFIED

Two observations, and the second is the point of this section.

First, on enforceability: ISO is a voluntary membership body. A Technical Specification obliges no one, no regulator stands behind it, and no manufacturer is required to run it. A measurement method spreads because practitioners find it worth running. Any framing in which a research institute seeks permission from such a document mistakes the nature of the document.

Second, on coverage: **the morphology whose energy signal is most informative about computation is the one excluded by name.** Section 4 established that an autonomous mobile robot spends more power on its cyber subsystem than on its motors. Section 5 and section 6 concern a humanoid, which the method does not mention. The bodies Physical AI actually deploys are precisely the bodies with no published method.

> **The Institute's position.** This is an opening rather than an obstacle. A method that is open, free to run, and executable rather than merely described is more adoptable than a paywalled document, and the fields that make it reproducible are already stated in section 7: the plane, the instrument's noise floor against a known load, the null control, and the residual so the floor can be reproduced rather than accepted. Publishing that freely, with the instrument alongside it, is the Institute's contribution here. A measurement method held behind a paywall is enclosure of something that ought to be common, and the answer to enclosure is to publish the free version people actually use.

9. What is this report careful not to claim?

Three of the six hypotheses are not settled by this evidence, and one of them was weakened by finding that somebody else got there first.

9.1 Custody, when the machine can be taken (H4)

Reducing the integrity of a metered system to custody of its measurement reference is a strong move, and it imports a mature body of practice: legal metrology has handled trusted measurement in hostile commerce for a very long time, through type approval, physical seals, traceable calibration chains and periodic recertification. That practice assumes the instrument returns for inspection. A deployed body may be stolen, altered in the field and returned to service, or simply leave the jurisdiction whose approval regime gives the seal meaning. What an energy reference must do to survive loss of custody is an open engineering question and this report does not answer it. Candidate directions include bounding attestation freshness by recertification interval, references that destroy rather than reveal on tamper, and binding a signature to a physical characteristic of the drivetrain rather than to a seal.

> **Update since v1, on the third candidate.** The drivetrain-signature route has since been characterised in this lab's working analysis and it is bounded more tightly than "candidate" suggests. Unit-to-unit manufacturing spread in the copper-loss coefficient is roughly **21.5 %** from published motor tolerances, against the **2.94 %** resolution floor of §5, so the signal exists. But of the four per-joint coefficients only that one is individually identifiable, and at a realistic resolution it carries about **2.8 bits**. Identity therefore scales with joint count: roughly **13.5 bits** for the seven-joint arm of §5, and about **56**

bits for a full humanoid. Two consequences follow and both are restrictive. Uncompensated thermal drift of **0.633 %/°C** exceeds the entire manufacturing spread by a **34 °C** rise, which a motor passes through in normal use, so the signature dissolves without a thermometer. And at these entropy levels the mechanism can support *provenance verification*, testing whether a unit is the one enrolled, but **not cryptographic key derivation**, which no realistic joint count reaches. Those are different primitives and this report does not conflate them. The remaining open quantity is how far wear moves the coefficient over a service life.

9.2 A receipt is not evidence (H5)

This one holds by construction and is stated plainly because this Institute owns the artifact it constrains. An energy receipt computed by a host is an accounting record of what that host believes it did [6]. If the host is compromised, the receipt is the compromised host's account of itself, and signing it changes its provenance rather than its truth. A receipt becomes an integrity claim only when bound to a measurement the host cannot author. Nothing in this report establishes that binding; it establishes only that the binding is the load-bearing part.

9.3 Task disclosure, where prior art already exists (H6)

An actuation trace carries payload, trajectory, route and duty cycle at a resolution requiring no specialist access, and the operational exposure of that is arguably larger than the exposure of any model parameter. This review initially treated that as an unclaimed observation. It is not, and the prior art is closer than expected on both halves.

The inference is established, on production machinery specifically. Gebbe measures only the aggregate three-phase current and voltage of a production machine at 10 kHz and disaggregates it into its constituent electric loads, recovering **the active power demand and operation schedule of individual loads** to an accuracy of **85 % or better** in two case studies, a thermoform machine and a milling machine [8]. The load taxonomy behind the method comes from an analysis of **151 electric loads** in manufacturing machines: fixed-speed motors 38 %, variable-speed-drive motors 26 %, Joule heating 15 %, rectifiers 13 %. Nothing in that work is framed as disclosure; its stated purpose is energy efficiency. But the capability it demonstrates is precisely the one a machine's operator would not want an observer to have.

The privacy consequence is also established, and it already names the industrial case. Giaconi, Gündüz and Poor survey privacy-aware smart metering and state directly that beyond residential users, metering privacy is critical for businesses "*e.g., factories and data centers, as their power consumption profile may reveal sensitive information about the state of their businesses to their competitors*" [9]. They classify the countermeasures into two families: **data manipulation** (obfuscation by added noise, aggregation, anonymisation, down-sampling, including differential privacy applied to the meter stream) and **user demand shaping**, which modifies the actual consumption using physical resources such as rechargeable batteries and local generation. Both families are mature and neither was developed for robots.

The contribution therefore reduces to a transfer: an established inference applied to an embodied platform, where the disclosed quantity is a physical task rather than a household activity, and where the existing countermeasures appear not to have been applied. **The novelty claim is withdrawn.** The discipline is the one that governs the whole cluster: claim the transfer, never the method.

> **A note on the second countermeasure family, because it lands differently on a body.** Demand shaping hides a load by changing what the system actually draws, using a battery to decouple consumption from activity. On an embodied system that option is constrained in a way it is not in a building: the actuation term is doing mechanical work on the world, and the work cannot be rescheduled without changing what the machine did. A robot cannot shift its lifting to off-peak. So of the two families, the manipulation family transfers directly and the shaping family transfers only to the compute term. This review did not locate a treatment of that asymmetry.

10. How did each hypothesis resolve?

#	Position on the trajectory	Binding constraint	Threshold that would move it
H1	Supported. Plane dependence measured at both ends: ~1.9× across four planes on an industrial plant [6], and a humanoid arm rail 15 W	Engineering implementation	Per-rail energy measurement as a standard build option rather than a

	below whole-robot draw while excluding most leg and computer loads [2]		laboratory modification
H2	Refuted as written, survives weakened. Physics-structured, parameters identified; 10.1 % mechanical [2]	AI algorithms and engineering implementation	A loss model whose coefficients are predicted from motor and gearbox datasheets rather than fitted per unit
H3	Supported and quantified. 10.2 % of trials contaminated by cross-subsystem energy transfer [2]	Engineering implementation	A second, non-electrical modality instrumented by default, so closure is checked two ways
H4	Open. No candidate mechanism evaluated here	Engineering implementation, and regulation for the approval regime	A reference that survives adversarial physical possession, with a stated freshness bound
H5	Holds by construction	Engineering implementation	A measurement path the host cannot author, bound to device identity
H6	Prior art located; a port. Novelty withdrawn [8,9]	None technical; the gap is application	An existing NILM countermeasure evaluated on an embodied platform's actuation trace

Table 3. Four of six bind on engineering implementation. That is the finding at the level of the programme: the physics carries and the instrumentation is missing.

11. What would sharpen this?

Three additions would each move this map further than any amount of re-reading the sources it already has.

Widen the evidence base. It is three primary documents. That is small, and it is stated in the badges rather than obscured by breadth. The two robot studies are single-platform: one research AMR with a Jetson board [1], one Unitree G1 arm [2]. Neither result is established as general across their morphology classes, and this report does not treat them as such.

Reference 2's authors record three limitations that propagate into anything built on their residual: motor winding resistance drifts with temperature and is uncorrected over long runs; harmonic drives are not back-drivable, so braking efficiency is asymmetric while the model uses one coefficient for both directions; and where two joints follow identical velocity profiles, viscous and interaction terms are not separately identifiable.

Measure the split on a humanoid or a legged platform. This is the single highest-value figure anyone could add to this map, and a measured compute-versus-actuation split for such a platform was not located by this review. The searches run covered legged-robot power breakdowns, industrial-robot controller standby share, and humanoid compute budgets; what returned at that specificity were vendor and secondary sources, which are not cited here. Reference 1's split is for an AMR and should not be transferred to a humanoid without measurement.

Read the clause text of reference 3. It was not read for this report. Whether ISO/TS 25213 fixes a measurement point internally is therefore unknown to this report, and is treated as unknown rather than assumed in either direction. Since the specification obliges no one, this is a gap in our knowledge of their document rather than a gap in the argument.

References 8 and 9 were read in full for this revision, and the H6 discussion now reports their measured content rather than their existence. A third located source on NILM applied to induction motors in textile manufacturing was not read and has been removed from the reference list rather than carried unread.

12. Conclusions

Energy is a real substrate for a machine to demonstrate its own behaviour, and its reach on an embodied system is set by the body rather than by the architecture. Three numbers carry the report. Under autonomy, an AMR's motors draw 16.6 percent of system power against the GPU's 37.3 percent, so computation is the larger term and the masking intuition is backwards. A physics-structured model of a humanoid arm predicts electrical power to 3.58 W at unseen operating points, which is the detection floor, while only 10.1 percent of that power is mechanical work, so the model's structure is physics and its magnitudes are fitted. And 10.2 percent of trials in that same study had to be discarded because energy crossed a subsystem boundary the meter could not resolve, a confound found with an inertial sensor rather than an ammeter and present without any adversary at all.

Four of the six hypotheses bind on engineering implementation, which is the programme-level result: the physics carries, and the instrumentation is absent. Commodity platforms expose a battery gauge, not per-rail energy. The one published measurement method covers six-axis industrial arms, obliges nobody, and excludes by name the morphology where the energy signal is most informative about computation. The bodies Physical AI actually deploys have no method, which is the ground this Institute intends to occupy, freely and in the open, with the instrument alongside the method so that both can be run rather than merely cited.

Two claims made in this Institute's own drafting were refuted by the evidence and one novelty claim was withdrawn. They are reported here at the same size as the findings, because a report that lists only what survived is not reporting.

13. The forcing function

Every bound in this report is a specification for the next piece of work, and it is worth stating them that way round, because a bound read as a verdict stops enquiry and a bound read as a threshold starts it.

What is bounded	The physics that sets it	The engineering change that moves it	What becomes possible when it moves

Compute integrity on a whole-system meter	The residual of a power prediction against the noise of everything else the body is doing	**Per-rail energy sensing as a standard build option**, a component decision, not a discovery	A machine that reports what it computed, on any body, without a laboratory
The expected value of an embodied load	Only 10.1 % of electrical power becomes mechanical work; the rest is loss with fitted coefficients	Loss coefficients predicted from motor and gearbox data rather than identified per unit	A null that needs no enrolment, so a machine can be checked the day it is built
Closure of the energy balance	Energy crosses subsystem boundaries as mechanics, not as error	A second, non-electrical modality instrumented by default	An audit equation that closes on a body, which is what makes an unattended machine accountable at all

Table 4. None of the three requires new science. Each is an instrumentation decision that has not yet been made, which is the most tractable kind of obstacle there is.

The through-line: measuring energy on a body was *not possible* in any useful sense a decade ago, is *demonstrably tractable* today at the planes and residuals reported here, and becomes **ordinary** the moment per-rail sensing is a checkbox on a build sheet. That is the shape of the transition this Institute exists to shorten.

References

1. L. Liu, W. Shi and K. G. Shin, "Power-Efficient Autonomous Mobile Robots," arXiv:2511.20467v1, 25 November 2025. READ IN FULL
2. N. N. Deniz, S. Vega, S. Parsons and F. A. Auat Cheein, "Identification of a Physics-Based Electrical Power Consumption Model for the Unitree G1 Humanoid Arm," arXiv:2606.15915v1, 14 June 2026. READ IN FULL
3. 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
4. International Federation of Robotics, "ABB Robotics leads global effort to standardize measurement of industrial robots' energy consumption," 2026. VERIFIED
5. NVIDIA, Jetson AGX Orin series module specifications: 15–60 W configurable power, up to 275 TOPS. VENDOR SPECIFICATION

6. Institute for Physical AI @ JBI, TR-2026-06, *MathGround: joules, not tokens*, and PAI-280, *Measuring Energy in Physical AI: Where You Put the Meter*. INSTITUTE PUBLICATIONS
7. Institute for Physical AI @ JBI, TR-2026-07, *Provable by construction*. INSTITUTE PUBLICATION
8. C. Gebbe, "Nonintrusive Load Monitoring for Machines used in Manufacturing," arXiv:1912.01500v1, 1 December 2019. Fraunhofer Research Institution for Casting, Composite and Processing Technology (IGCV), Augsburg. READ IN FULL
9. G. Giaconi, D. Gündüz and H. V. Poor, "Privacy-Aware Smart Metering: Progress and Challenges," arXiv:1802.01166v4, 24 May 2018. Imperial College London and Princeton University. READ IN FULL

▶ Run the instrument behind this report