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

SECURITY · PROVENANCE

Technical Report TR-2026-37
Research / Position · Preprint v1
13 August 2026

SECURITY · PROVENANCE

Can a Machine Prove It Is Itself?

A drivetrain carries a measurable manufacturing signature. This report prices it in bits, and the price decides what it can be used for.

The Charlot Lab · Institute for Physical AI @ John Bailey Institute
Companion to TR-2026-36, which states this question and does not answer it.

4 questions resolved without hardware | 2.8 bits per coefficient | the decisive question was posed wrong | 1 of our own extrapolations withdrawn

A body that can be stolen, altered or substituted needs a way to demonstrate it is the body that was certified, and software cannot supply one because software is what an adversary replaces. Manufacturing variance is the classical answer: silicon physically unclonable functions turn fabrication accidents into identity. This report asks whether a drivetrain does the same, and answers entirely from published coefficients and computation. It does. Vendor tolerances put unit-to-unit spread in the copper-loss coefficient at about 21.5 percent against a model resolution floor of 2.94 percent, so the signal clears the noise by roughly seven times [1,2]. But three constraints bound it hard. Only one of the four per-joint coefficients is individually identifiable, at a standard error of 0.2 to 1.0 percent, while the others sit at 36 to 54 percent and cannot carry an identity at all. Temperature drift compounds, resistance rising and magnet strength falling both push the same term the same way at 0.633 percent per degree, so the entire manufacturing spread is consumed by a 34 degree rise, which a motor passes through in ordinary use, unless winding temperature is measured. And the binding constraint turns out to be neither of those: it is **entropy**. At realistic resolution one coefficient carries about 2.8 bits, and since only one coefficient per joint is identifiable, identity scales with joint count, roughly 13.5 bits for a seven-joint arm and 56 for a full humanoid. That is enough for **provenance verification**, testing whether a unit is the one enrolled, and it is not enough for cryptographic key derivation, which no realistic joint count reaches. The wear question that opened this work was the wrong question: wear is one contribution to a bit-error budget that a fuzzy extractor is designed to

absorb [7], and the single measured wear datum available refutes the linear extrapolation this report attempted from it.

1. What is the question, and where does it come from?

TR-2026-36 reduced the integrity of a metered system to custody of its measurement reference, and then observed that the reduction imports an assumption which embodiment breaks: legal metrology's apparatus of type approval, seals and periodic recertification assumes the instrument comes back for inspection. A deployed body may be stolen, altered in the field and returned to service, or simply leave the jurisdiction whose approval gives the seal meaning. That report named the question and declined to answer it.

This one answers as much of it as can be answered without hardware, which turned out to be all four of the questions it opened with. The result is more restrictive and more useful than expected, and the most important finding is that the question this work began with was not the binding one.

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

Every figure carries a grade. **Verified** means a primary document was retrieved and the number read from it. **Vendor** means a manufacturer's published specification. **Computed** means our arithmetic over graded inputs, and the arithmetic is shown. **Not located** means this review searched for a quantity and did not find it, and the searches are named.

No hardware was used and none was needed. That is itself a finding: the first draft of this work opened with a request for five identical actuators, and every question it was meant to settle was already answerable from published specifications and computation.

3. What is already known, and what does this add?

Field	What it owns	Granularity

Silicon PUF	Per-unit identity from manufacturing variance; the canonical form of this idea	individual unit, in silicon
Motor current signature analysis	That individual motors have learnable baselines, and systems that learn them per motor	individual unit, framed as condition monitoring
EMI fingerprinting	Counterfeit detection from electromagnetic emissions	make, model, configuration; patent-dense

Table 1. Three fields own parts of this ground. The physics we lean on, that nominally identical motors differ enough to be modelled individually, is established by the second of them. Our contribution cannot be, and is not, "motors differ".

This report claims no fingerprinting method. It states a physical question, prices the answer in bits, and publishes both openly with the arithmetic attached. That is a deliberate position: a measurement method held behind a patent or a paywall is enclosure of something that ought to be common, and the answer to enclosure is to publish the free version.

What this review did not locate is a published test of whether a physics-structured energy model's identified parameters separate individual drivetrain units of the same make and model. Searches run were on physically unclonable functions combined with motor and drivetrain manufacturing variance, motor current signature analysis and unit-to-unit variation, and device authentication and counterfeit detection from motor electrical signatures. That is a statement about this review, not about the literature.

4. Is there a signal? Yes, and the number is on a datasheet

The model is the one identified in TR-2026-36 §5 for a humanoid arm [2]: per joint, a mechanical term, a baseline-corrected copper-loss term, and Coulomb and viscous friction terms. The copper-loss coefficient is $b = R / (n \cdot \eta_g \cdot K_m)^2$, so it carries phase resistance once and the torque constant twice.

Vendors publish unit-to-unit tolerance because design engineers need it to size a system that works with any unit off the line. For brushless DC motors, torque constant is quoted at **±10 %**, of which magnet remanence contributes roughly 2.5 % and phase setting 7.5 %, and phase resistance at ±8 % [1]. VENDOR

Route	Implied spread in b	Against the 2.94 % floor
Worst case, linear	28.0 %	9.5x
In quadrature	21.5 %	7.3x
Combined datasheet figure alone	18 %	6.1x

Table 2. The signal clears the model's own resolution floor by roughly an order of magnitude. COMPUTED FROM VENDOR TOLERANCES

> **The methodological point is worth more than the number.** This quantity was originally scoped as an experiment requiring five identical actuators. It is a specification. Before designing a measurement, check whether the manufacturer has already published the variance, because tolerance *is* a statement about how much units differ.

5. Is it identifiable? Only one coefficient, and excitation decides

A fingerprint must be a property of the hardware and not an artifact of the fitting procedure. If two coefficient values explain the data equally well, whichever one the optimiser returned is an accident, and it will be a different accident on the next fit. So identifiability precedes distinctiveness.

We built the per-joint regressor on a representative single-joint plant, using the source's own trajectory design, and computed each parameter's standard error at the model's 1.07 W identification residual.

Excitation set	cond($\Phi^T\Phi$)	b copper	a mech	d viscous
1 travel, 1 posture	3.4e17	321 %	–	3811 %
3 travels	2.4e07	1.0 %	792 %	121 %
3 travels, 4 postures	5.6e05	0.3 %	54 %	54 %

4 travels, 6 postures	1.1e06	0.2 %	37 %	36 %

Table 3. Standard error as a fraction of each coefficient's value. Only b resolves tightly enough to be a property rather than a choice. COMPUTED

Two readings. First, **only the copper-loss coefficient can carry an identity**: 0.2 to 1.0 percent against the 21.5 percent it must resolve, while the mechanical and viscous terms sit at 36 to 54 percent and would contribute mostly noise. That it is the same coefficient the datasheet tolerances point at is not a coincidence, b contains both quantities that vary most between units, but the two arguments are independent and they agree.

Second, **excitation is decisive**. Adding two travel distances moved the conditioning by ten orders of magnitude. Identifiability is a property of the experiment as much as of the model, and a naive sweep set will not identify these parameters at any instrument quality.

> **A result of ours that was wrong, and how it was caught.** The first run of this analysis reported that no coefficient was identifiable. It also inverted the source's own conditioning finding, which reports per-sample fitting as degenerate and trajectory-averaging as sound; our run showed the reverse. That contradiction was the signal. Five near-collinear rows from a one-parameter trajectory family cannot identify four parameters, the instrument was broken, not the method. A red result that disagrees with a published one about the same diagnostic is a bug until proven otherwise.

6. Does it survive temperature? Not without a thermometer

Copper resistivity rises with temperature at **+0.393 %/°C**. NdFeB remanence, and therefore the torque constant, falls at **−0.08 to −0.12 %/°C** [3]. Since b carries resistance in the numerator and the torque constant squared in the denominator, **a rising numerator and a falling denominator push the same way and the effects compound** rather than cancelling: **+0.633 %/°C**. COMPUTED FROM PUBLISHED COEFFICIENTS

Temperature rise	Drift in b	Against the 21.5 % spread
20 °C	12.7 %	0.59x
34 °C	**21.5 %**	**1.00x, crossover**

60 °C	38.0 %	1.77x

Table 4. Beyond roughly 34 °C the drift exceeds the entire difference between units, so a warm unit reads as a different cold one. A motor passes through that between cold start and steady running.

The remedy is cheap and it is sufficient. Correcting to a reference temperature leaves a residual set by how well winding temperature is known: **±5 °C leaves 3.2 %**, about a seventh of the signal; ±1 °C leaves 0.6 %. **The fix is a thermometer, not a better ammeter.**

7. What is the binding constraint?

The three sections above answer the questions this work opened with, and none of them is the constraint that decides the outcome. A fuzzy extractor, the standard construction for turning a noisy physical source into a stable identity, with an enrollment phase that computes public helper data through an error-correcting code and a reconstruction phase that recovers the identity from a noisy re-reading [7], is denominated in **bits**. For a continuous coefficient the usable bits are approximately $\log_2$(spread / resolution).

Condition	Effective resolution	Bits per coefficient
Model floor only	2.94 %	2.87
Thermometer to ±5 °C	3.16 %	2.76
No thermometer, 40 °C swing	25.3 %	−0.24, none

Table 5. The last row confirms §6 from an independent information-theoretic direction: uncompensated, the source carries negative usable entropy. Two arguments, one conclusion. COMPUTED

Because only one coefficient per joint is identifiable, **identity scales with joint count**, at about 2.76 bits each, less the entropy the public helper data leaks.

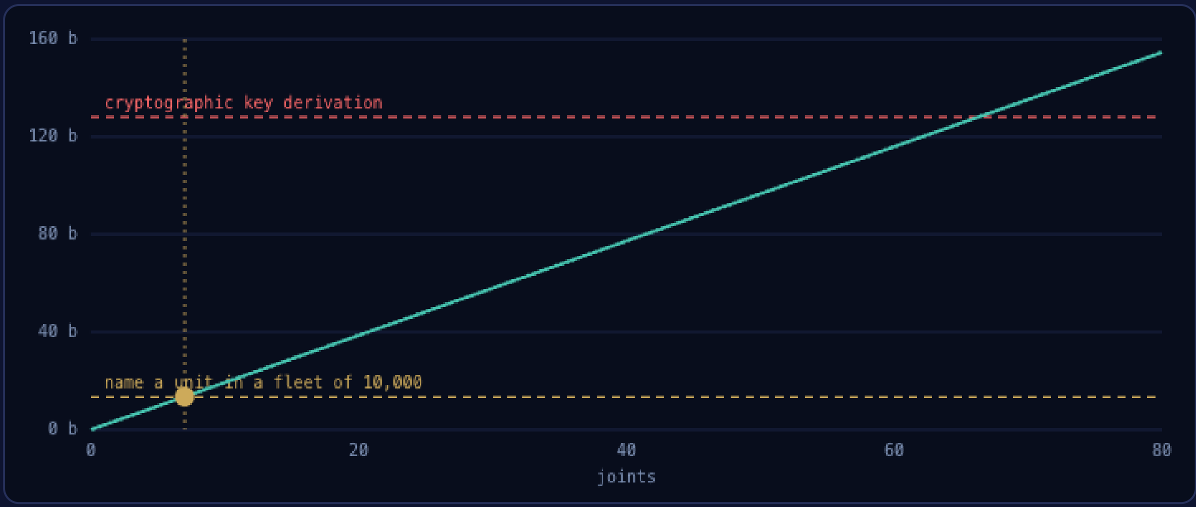

Joints **7** Helper leak **30 %**

7 joints → 13.5 bits · false accept about 1 in 11,780 · can name a unit in a fleet of 10,000; still far below key derivation

Figure 1. Drag the joint count. A seven-joint arm reaches about 13.5 bits, a full humanoid about 56, and the 128-bit line needed for key derivation is not reached below roughly 67 joints. Predict before dragging: most readers expect a robot to be able to generate a key from its own body. COMPUTED

8. Which of the three primitives is defensible?

Entropy separates three security primitives that are easy to conflate, and the PUF literature is written mostly about the third of them.

Primitive	What it asks	Cost	Verdict at these levels
Verification	Is this the unit I enrolled?	one hypothesis	supported, an arm's 13.5 bits is a ~1-in-11,600 false accept
Identification	Which of N units is this?	$\log_2(N)$ bits before anything else	marginal, 13.3 bits merely to name a unit among ten thousand
Key derivation	Generate a secret from the body	~128 bits	not supported, needs ~67 joints

Table 6. The claim this report makes is the first row and only the first row.

> **The defensible claim is provenance verification, not key generation.** A drivetrain cannot be a cryptographic root of trust at any realistic joint count. It can, on the arithmetic above, tell you whether the actuator in front of you is the one that was enrolled, which is precisely the question a stolen, altered or substituted body raises, and it needs no cryptographic strength to be useful.

9. Why was the question this work opened with the wrong one?

This report began by treating wear as the decisive experiment: if the coefficient drifts as a drivetrain ages, the mechanism dies. That framing was binary where the real answer is a budget. In a fuzzy-extractor construction, drift does not have to be zero, it has to sit inside the correction capacity, alongside every other noise contribution [7].

What was not located: a published coefficient for how a drivetrain's copper or friction terms drift with accumulated duty hours. Searches covered motor parameter drift and winding-resistance aging, harmonic-drive efficiency degradation over operating hours, and accelerated life testing of servo actuators. The direction is uncontroversial in the maintenance literature (resistance rises with thermal aging, gear friction rises with wear and lubricant degradation) but no magnitude usable here was found.

What was located is a single measured datum, and it is only weakly analogous: in a 57-hour gait-like life test of a 3D-printed open-source actuator, positive-work efficiency fell by **3.5 %**, levelled off, and then **rebounded to near its original value**, which the authors attribute to transmission dust reaching saturation over an underlying break-in process [6]. `VERIFIED, WEAK ANALOGY`

Since b carries gear efficiency squared in its denominator, a 3.5 percent efficiency loss implies a **7 percent** drift in the coefficient, a third of the manufacturing spread, and equivalent to an 11 °C temperature change on the thermal axis. Combined with the residual from a ±5 °C thermometer, the total sits at 10.2 percent, inside the 21.5 percent budget but consuming **47 percent of it**, and costing **1.68 of the 2.76 bits** per coefficient.

An extrapolation of ours, withdrawn on its own source. The obvious next step is to convert 3.5 % over 57 hours into a rate and project it. We did, and it says the identity is gone by 500 hours, which contradicts the very observation the datum comes from, since the source reports the efficiency *recovering*. Wear here is not monotonic, a constant-rate model is unjustified, and the absurd projection is the tell rather than a result. **The linear extrapolation is withdrawn.** What remains is that a long-duration measurement on a representative drivetrain is the one thing in this report that genuinely requires hardware and time.

And the mitigation is already standard practice in the field this work started from. If the identity is **re-enrolled periodically**, wear drift only has to stay inside the correction capacity *within the re-enrollment interval*, not across a service life. That is the freshness bound TR-2026-36 §9.1 sketched from first principles, and it is what legal metrology already does under the name recertification. The physics converges on the operational practice.

10. How did each hypothesis resolve?

Question	Answer	Binding constraint	Threshold that would move it
Is unit variance above the floor?	yes, 7.3x	none	—
Is any coefficient identifiable?	yes, b only	experiment design	an excitation set covering ≥3 travels and ≥4 postures
Does it survive temperature?	no uncompensated, yes at ±5 °C	engineering implementation	winding-temperature sensing as a build option
Does it survive wear?	open; one weak datum	material science	a long-duration efficiency curve on a representative drivetrain
Is there enough of it?	for verification, not for keys	**entropy, the binding one**	no instrument improves it; only joint count does

Table 7. Four of five resolved from published coefficients and computation. The fifth needs hardware, and it is not the one this work expected.

11. What would sharpen this?

Four measurements, each of them a bench session rather than a research programme.

Fit a batch of real units. The §5 analysis simulates the regressor's structure on a representative plant, which establishes which parameter *can* carry an identity given the model form and a stated residual. Fitting a batch of physical actuators is what would turn that from inferential to established, and it needs a test bench rather than a new method.

Characterise wear on a harmonic drive. The §9 datum comes from a 3D-printed actuator whose dominant wear mechanism is transmission dust, so it sizes the term. A harmonic-drive wear study would characterise it, and the term is the one most likely to move the result.

Compute the entropy rigorously. The estimate treats $\log_2$(spread / resolution) as usable bits with a nominal 30 percent helper-data leak, both first-order. Computing the min-entropy of the actual coefficient distribution against the exact leakage of a chosen code would replace both, and would most likely return a smaller number rather than a larger one, which is worth knowing before anyone builds on it.

Measure the correlation between joints. Coefficients on one machine are treated as independent here. Shared supply voltage, a common thermal environment and correlated manufacturing batches all pull against that, so the joint-count scaling of §7 is an upper bound until somebody measures how far.

12. Conclusions

A drivetrain does carry a manufacturing signature large enough to see: about 21.5 percent unit-to-unit spread in the copper-loss coefficient against a 2.94 percent resolution floor. Only that one coefficient of four is identifiable well enough to use. It dissolves under ordinary thermal operation unless winding temperature is measured, and measuring it to ±5 °C is sufficient. And the constraint that actually decides the outcome is none of those: at about 2.8 bits per coefficient, a seven-joint arm carries roughly 13.5 bits and a full humanoid about 56.

That supports **provenance verification** and not cryptographic key derivation, and the distinction is the report's main contribution, because the literature this construction comes from is written mostly about the primitive a drivetrain cannot support.

Two of this report's own positions did not survive its own analysis: the wear question it opened with was not the binding constraint, and the linear extrapolation it attempted from the single available wear datum is contradicted by that datum's own source. Both are reported at the size they were claimed.

13. The forcing function

This report's conclusion is restrictive, and restriction is useful only if it names what would lift it.

What is bounded	The physics that sets it	The engineering change that moves it	What becomes possible when it moves
~2.8 bits per coefficient	Manufacturing spread divided by model resolution	Either tighter resolution or **more identifiable coefficients per joint**, the second is an experiment-design problem and §5 shows the excitation set moves conditioning by ten orders of magnitude	Fleet-scale identification rather than pairwise verification
Identity dissolves at 34 °C	Copper resistivity rises while magnet remanence falls, and the two compound in the same coefficient	**Winding-temperature sensing to ±5 °C**, which is a thermistor	A signature that survives its own duty cycle, which is the difference between a laboratory result and a deployable one

Wear over a service life	Friction and resistance drift with hours; no published coefficient located	Passive current logging on a machine already running, plus periodic re-enrolment	Provenance that holds across a machine's whole life rather than its first week

Table 8. The first row is the one that decides how far this goes, and it is an experiment-design question rather than a materials one.

Stated as a trajectory rather than a verdict: a drivetrain carrying a checkable identity was *not a thing anyone could claim* before the manufacturing spread and the model residual were put in the same units; it is *demonstrably sufficient for verification* today; and it becomes **sufficient for identification across a fleet** once the number of identifiable coefficients per joint rises. That is a solvable problem in excitation design, and it is the next thing to build.

References

1. Published brushless DC motor unit-to-unit tolerances: torque constant ±10 % (magnet remanence ≈2.5 %, phase setting ≈7.5 %), phase resistance ±8 %, combined electrical up to ±18 %. VENDOR SPECIFICATION
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. Published NdFeB reversible temperature coefficients of remanence, −0.08 to −0.12 %/°C; torque constant tracks flux. VENDOR AND MATERIALS LITERATURE
4. Institute for Physical AI @ JBI, TR-2026-36, *Energy Observability in Embodied Systems*. INSTITUTE PUBLICATION
5. Institute for Physical AI @ JBI, PAI-290, *Energy as Evidence*, Module 3. INSTITUTE COURSE
6. Design and Characterization of 3D Printed, Open-Source Actuators for Legged Locomotion, arXiv:2202.12395. 57-hour life test; positive-work efficiency fell 3.5 %, levelled, then rebounded. VERIFIED FOR THE CITED FIGURE
7. A. Schaller, B. Škorić and S. Katzenbeisser, "On the Systematic Drift of Physically Unclonable Functions Due to Aging," in *Proceedings of the 5th International Workshop on Trustworthy Embedded Devices (TrustED 2015)*, Denver, Colorado, ACM, pp. 15–20. ISBN 978-1-4503-3828-8, doi:10.1145/2808414.2808417. BIBLIOGRAPHIC RECORD VERIFIED; FULL TEXT NOT READ

▶ Run the instruments behind this report

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

TR–2026–37 · Preprint v1 · 13 August 2026