An Open Letter to the Quantum-Computing, Thermodynamic-AI, and AI-Evaluation Communities: Thermodynamic, Algebraic, and Psychometric Bounds on Hybrid Quantum–Photonic and Neuro-Symbolic General Intelligence

Peter De Ceuster*

*Quantum & Thermodynamic Computing Working Note: Sagawa–Ueda Guardrail Accounting, Lambert-W Context Scaling, and Behavioral Meta-Cognition,
(peterdeceuster5@gmail.com)

November 19, 2026

Abstract

This document is written as an open letter to researchers in topological quantum computing (including teams at Microsoft Development Center Copenhagen and NIST), thermodynamic AI hardware, formal verification, and AI evaluation. Across sixteen prior works and a subsequent Majorana–photonic core synthesis [1–17], we proposed sub-Landauer optomagnetic logic, interferometric Majorana sector-gates, cohomological sheaf guardrails, closed-form context-abstraction limits, the AISA-T meta-cognition benchmark, and game-theoretic dissipation models. A rigorous physical, algebraic, and psychometric audit uncovers critical gaps: an unaccounted thermodynamic cost of dissipative sheaf/radiative guardrails, a ten-to-thirteen-order-of-magnitude drive and Carnot omission in optomagnetic switching, a dropped scale factor in the context boundary reducing 1.21×10^6 tokens to 7.90×10^4, non-trace-preserving Kraus operators paired with incompatible JVM/n8n latencies, Jacobi-violating neural GENERIC brackets, confabulatory AISA-T self-reports, and misstated 2021 World Championship chess dynamics.

Here we establish a self-consistent theoretical and benchmark synthesis. First, we prove the *Sagawa–Ueda / Landauer Guardrail Theorem* quantifying the exact wall-plug energy bill of cohomological filtering, magic-state distillation, and cryogenic refrigeration alongside Margolus–Levitin and Sidis open-system bounds. Second, we derive the exact Lambert-W context-abstraction boundary and repair the reduced density, commutator, Chernoff, Kraus, and diffeomorphism-pullback GENERIC equations. Third, we reform AISA-T into a proper-scoring behavioral protocol (AISA-T v2.0) and formulate a 50-move-clock absorbing Quantal Response Equilibrium chess Markov model.

What follows is a technical formulation of these thermodynamic, algebraic, and psychometric resolutions, together with falsifiable laboratory and software evaluation protocols. The discussion is organized in six points.

(1) The Sagawa–Ueda / Landauer Guardrail Theorem and thermodynamic energy accounting. A central unresolved tension across [11, 12, 17] is the *reversible-core versus dissipative-verifier paradox*: while interior unitary Majorana braiding incurs zero logical erasure cost, the cohomological guardrails of [12, 17]—sheaf heat diffusion $ds/d\tau = -\gamma_{\text{diff}}\Delta_{\mathcal{F}}s$, radiative damping Γ_{mn}, upstream rejection, and parity readout reset—are strictly logically irreversible contractions from M candidate hypotheses to K consistent sections in $H^0(X;\mathcal{F})$. By Landauer's principle, Sagawa–Ueda feedback thermodynamics, and Carnot refrigeration from room temperature T_{env} to cryogenic operation T_{op} ($T_{\text{env}}/\eta_{\text{Carnot}}T_{\text{op}} \sim 10^3$–$1.5\times10^4$), every cycle satisfies the wall-plug bound alongside the Margolus–Levitin quantum speed limit in Eq. (1):

$$E_{\text{wall-plug}} \geq \frac{T_{\text{env}}}{\eta_{\text{Carnot}}T_{\text{op}}}\left[k_B T_{\text{op}}\ln\frac{M}{K} + N_{\text{meas}}\,k_B T_{\text{op}}\,\Delta H_{\text{reset}} + E_{\text{magic}}(N_T,\epsilon) + \alpha_{\text{loss}}E_{\text{drive}}\right], \qquad E_{\text{op}} \geq \frac{\pi\hbar}{2\tau_{\text{op}}}, \tag{1}$$

Applying Eq. (1) resolves three severe thermodynamic discrepancies in [10, 11, 16]: (i) in [11], comparing Q_{irr} at $T_{\text{op}} \lesssim 5\,\text{K}$ against $k_B T_{300\text{K}}\ln 2$ underestimated the true operating entropy ratio by $60\times$ (at $5\,\text{K}$, $4.2\times10^{-22}\,\text{J} = 8.8\,k_B T_{5\text{K}}\ln 2$, super-Landauer before Carnot overhead!), each $250\,\text{meV}$ drive photon already equals $13.95\,k_B T_{300\text{K}}\ln 2$, the drive pulse carries 10^{10}–$10^{13}\,k_B T_{300\text{K}}\ln 2$, parameters give $R_{\text{mix}} = 16.50$ (not 1.65), $\Delta_{\min} = 0.414\,\text{meV}$ (not $2.5\,\text{meV}$), $\tau_p^* = 23.79\,\text{ps}$ (not $12\,\text{ps}$), $Q_{\text{irr}}^{\min} = 1.67\times10^{-23}\,\text{J}$ (not $4.2\times10^{-22}\,\text{J}$), Eq. (29) omits a logarithmic factor ≈ 15.2, $b_{\text{elec}} + b_{\text{mix}} = 14.8\,\text{meV} < 25\,\text{meV}$ fails to invert the gap, Eq. (9) Magnus sign and Eq. (11) factor i (anti-Hermitian!) must be corrected as $-\frac{2e^2}{\hbar^3\Omega}$, Eq. (20) degree-mismatched $d(*F - \theta F - J_s) = 0$ must become $d(*F - \theta F) = J_s$, co-propagating TE/TM impedance matching $E_A B_B + E_B B_A$ is required so $\text{Im}(\mathbf{E}^* \cdot \mathbf{B}) \neq 0$, coherent Rabi π-pulses must replace irreversible SET latching for Toffoli gates, 2D Kibble–Zurek scaling replaces 0D Landau–Zener, and ordinary Bragg GVD replaces Lorentz-violating operators; (ii) in [10], human brains ($\sim 1.3\times10^{-13}\,\text{J/synOp}$) outperform NVIDIA B200 ($\sim 2.2\times10^{-13}\,\text{J/op}$) and Frontier ($\sim 1.9\times10^{-11}\,\text{J/FLOP}$) by ~ 1.7–$140\times$ rather than $10^{10}\times$, while $10^{20}\times$ compression below CMOS ($10^{-34}\,\text{J/op}$) violates Margolus–Levitin by 11–13 orders at 0.1–10 THz; and (iii) in [16], physical time reversal is $(t,S)\mapsto(-t,+S)$ (flipping both t and S preserves dS/dt), Prigogine's internal entropy production always satisfies $d_iS/dt \geq 0$ (local decrease occurs via negative external export $d_eS/dt < -d_iS/dt$), and 3D Maxwell–Boltzmann median-speed demon sorting yields $\Delta S_{\text{gas}} \approx -0.3251 N k_B$ rather than spatial compression $-Nk_B\ln 2$.

(2) Exact Lambert-W context boundary and quantum-control/GENERIC corrections. In [7], the context-abstraction limit dropped the scale factor τ inside the logarithm of Eqs. (R18) and (5), yielding a dimensionally invalid argument and $N^*_{\mathrm{eq}(5)} = 78{,}966$ instead of the exact root 1,211,096 for $N_{\mathrm{gate}} = 10^6, \tau = 10^5, A_0 = 100, \lambda = 1$ (while Eq. (R16) was an identical zero-step no-op and $\tau_{\mathrm{tunnel}} = 1/\gamma_{\mathrm{qp}}$ mixed seconds with tokens). Introducing the token rate $\tau_{\mathrm{tok}} = R_{\mathrm{tok}}\tau_{\mathrm{qpp}}$ gives the exact closed-form Lambert-W solution Eq. (2) (yielding 1,211,105 at leading order) and cleanly separates hardware poisoning from Transformer attention-entropy dilution:

$$N^* = \tau_{\mathrm{tok}}\, W_0\!\left(\frac{A_0}{\lambda} e^{N_{\mathrm{gate}}/\tau_{\mathrm{tok}}}\right) = N_{\mathrm{gate}} + \tau_{\mathrm{tok}} \ln\!\left(\frac{\tau_{\mathrm{tok}} A_0}{\lambda N_{\mathrm{gate}}}\right) - \tau_{\mathrm{tok}} \ln\!\left(1 + \frac{\tau_{\mathrm{tok}}}{N_{\mathrm{gate}}} \ln\frac{A_0}{\lambda}\right) + O\!\left(\frac{\tau_{\mathrm{tok}}^2}{N_{\mathrm{gate}}}\right), \tag{2}$$

Seven companion algebraic and architectural repairs across [1–9, 12] make the hybrid substrate rigorous:

- **Qubit geometry, reduced states, and commutators [1–3, 8]:** Pure single-qubit states live on $\mathbb{CP}^1 \cong S^2$ (not $\mathbb{C}^2$), projective readout obeys Holevo's ≤ 1 bit/qubit bound, and [2]'s reversed $|0\rangle/|1\rangle$ map is unified. Tracing out sector 3 of $|\Psi_2\rangle = \alpha|00\rangle_{12}|0\rangle_3 + \beta e^{i\phi}|11\rangle_{12}|1\rangle_3$ kills all cross-terms (Eq. 3, correcting Eq. 24 of [3]), while $[\hat{U}', \hat{V}'] = \hat{U}'\hat{V}' - \hat{V}'\hat{U}'$ is purely diagonal (Eq. 3, fixing the $[\hat{D}, \hat{G}] = \hat{U}' - \hat{V}'$ error shared by [3, 8]).
- **Clifford-only Majorana braiding and symmetric Chernoff exponent [4, 9]:** Majorana braiding and parity readout generate only the classically simulable *Clifford group*, so universal non-Clifford T gates require magic-state distillation overhead (contradicting Lemma 2 of [4]); Schur-concavity of $4t_L t_R|\sin\phi|$ maximizes contrast at $t_L = t_R$. In [9], doubling the branch separation and confusing KL divergence with symmetric Chernoff information overestimated κ_{rate} by $16\times$; Eq. (4) restores $\kappa_{\mathrm{Chernoff}}$ alongside $\Phi_0^{\mathrm{sc}} = h/(2e)$ and gauge-invariant boundary couplings.
- **Trace-preserving Kraus, CP SME, and cryo-CMOS control [5]:** Eq. (3.3) of [5] lost trace ϵ; Eq. (4) restores $\sum M_m^\dagger M_m = I$, while the stochastic master equation requires $+\mathcal{D}[\hat{c}]\rho_t dt$ for complete positivity. Because Repeat-Until-Success loops have unbounded worst-case latency and $\tau_{\mathrm{qpp}} \sim 1\,\mathrm{ms}$ forbids JVM GC pauses (1–100 ms), n8n latency (10–500 ms), and QMA-hard runtime diamond-norm proofs, feedforward must run on bounded-$K_{\max}$ cryo-CMOS/FPGA heralded hardware (< 500 ns).
- **Jacobi-exact GENERIC, Lean 4 interval predicates, and performative stability [6, 8, 12]:** Projecting arbitrary neural matrices $L_\theta = P_{dS}^{\perp}(A_\theta - A_\theta^T)(P_{dS}^{\perp})^T$ [12] generically violates the Jacobi identity (verified residual $1.532 \neq 0$) and breaks $\dot{E} = 0, \dot{S} \geq 0$ for $\varepsilon_0 > 0$. Pulling back a canonical Poisson tensor via an invertible diffeomorphism Φ_θ (Eq. 5) enforces Jacobi $\equiv 0$ identically, while Lean 4 verification requires computable interval propositions $\|\mathcal{F}_v s_v - \mathcal{F}_u s_u\| \leq \epsilon$ (rather than ill-typed Id = true) on a 2D Deductive×Empirical product lattice. In [8], Eq. (5) replaces unnormalized action KL with policy-weighted binary KL ≥ 0 under performative contraction $\|\partial p/\partial S\| < \mu_{\mathrm{cal}}$ benchmarked against Legg–Hutter $\Upsilon(\pi)$, and in [6], Theorem 1 requires upper-semicontinuous (p_M, J) and $\varepsilon < 1$ for `argmax` commit.

$$\rho_{12} = |\alpha|^2|00\rangle\langle 00| + |\beta|^2|11\rangle\langle 11|, \qquad [\hat{U}', \hat{V}'] = \left(e^{2i\varphi_1} - e^{2i\varphi_2}\right)\left(|110\rangle\langle 110| - |111\rangle\langle 111|\right), \tag{3}$$

$$\kappa_{\mathrm{Chernoff}} = \frac{(\Delta C_Q)^2}{4S_N}, \qquad M_m = \left(\sqrt{p(m|+1)}\,\Pi_+ + \sqrt{p(m|-1)}\,\Pi_-\right) U_{\mathrm{read}}, \qquad \sum_{m=\pm 1} M_m^\dagger M_m = I, \tag{4}$$

$$L_\theta(z) = D\Phi_\theta(z)\, L_0\, D\Phi_\theta(z)^T \ \ (\text{Jacobi} \equiv 0), \qquad E_S(t) = \mathbb{E}_{a\sim\pi_t}\left[p_t \ln\frac{p_t}{\hat{p}_t} + (1-p_t)\ln\frac{1-p_t}{1-\hat{p}_t}\right] \geq 0. \tag{5}$$

(3) Behavioral psychometric reform of AISA-T and bounded-rational chess dynamics. Auditing [13, 14] exposes both an arithmetic error (the ten self-scores of [13] sum to $269/300 = 89.67\%$, not 91%) and a severe construct-validity failure: top-rated GPT-5 answers fabricated nonexistent internal telemetry ("0.15 attention match ratio", "response composition layer", separation of text and state variables) while reporting $N = 1$ model, $N = 1$ unblinded rater, zero numeric human scores, and mismatched 10-point vs. 5-point scales. **AISA-T v2.0** replaces self-report introspection with proper-scoring behavioral probes—Brier score, Expected Calibration Error (ECE), selective-abstention AUROC, multi-hop recursion accuracy curves, and adversarial delimiter-injection tests—scored by ≥ 3 blinded raters (Krippendorff's $\alpha \geq 0.80$, ICC) across model and temperature sweeps.

Similarly, in [15], Game 6 of the 2021 World Chess Championship was misdescribed: Carlsen sacrificed *his* Queen for two Rooks on move 26, so Black kept a Queen to move 136 (not "queenless"); +0.15 pawns is 15 centipawns (not 0.15 centipawns); full moves were conflated with plies; deterministic minimax tanh collapses $V(s)$ to three values; Eq. (8) misused Shannon mutual information; Eq. (12) failed during move-expanding exchanges; and Proposition 7.3 ignored the 50-move draw clock ($c \in \{0, \ldots, 100\}$), which gives nonzero draw escape probability $\geq (1-\delta)^{100}$. Modeling two-sided Quantal Response Equilibrium (Boltzmann/Maia) error rates on a 50-move absorbing Markov chain with Stockfish 16+ MultiPV over Carlsen's 125-game unbeaten PGN corpus yields a rigorous empirical protocol.

(4) Quantitative benchmark summary. Table 1 contrasts earlier parameter claims across [1–17] with the corrected bounds.

Table 1. Quantitative corrections across hybrid AGI hardware, formal verification, psychometrics, and chess game theory.

Subsystem / Works	Earlier Claim / Defect	Corrected Bound / Protocol
Thermodynamics [10, 11, 12, 16, 17]	Sub-Landauer $10^{20}\times$; $T_{300\mathrm{K}}$ baseline; $\Delta S = -Nk_B \ln 2$	Eq. (1) Carnot + guardrail bill; $8.8\, k_B T_{5\mathrm{K}} \ln 2$; brain ~ 1.7–$140\times$ CMOS; $-0.3251 N k_B$
Context & Control [3–5, 7, 9]	$N^* = 78{,}966$; Clifford+T braid; $16\times$ κ; JVM/n8n RUS	Eq. (2) $N^* = 1{,}211{,}096$; magic-state T; $\kappa_{\mathrm{Chernoff}}$; FPGA $< 500\,\mathrm{ns}$
Neuro-Symbolic [6, 8, 12]	Jacobi residual 1.532; Lean `Id=true`; negative KL	Eq. (5) pullback L_θ (Jacobi $= 0$); interval Lean; binary KL + contraction
Evaluation & Chess [13–15]	91%; $N = 1$ unblinded; Game 6 "queenless"; no 50-move rule	89.67%; AISA-T v2.0 (Brier/ECE, $\alpha \geq 0.80$); 50-move QRE + Stockfish 16+

(5) Request to Quantum-Computing, Thermodynamic-AI, and Evaluation Communities. We request that topological-quantum groups (including MDCC) benchmark the wall-plug guardrail bound (Eq. 1) and bounded-$K_{\max}$ cryo-CMOS readout (Eq. 4) on InAs/Al parity experiments, verify the Jacobi-exact diffeomorphism GENERIC architecture (Eq. 5) on chaotic Vlasov–Maxwell rollouts, and execute the blinded multi-model AISA-T v2.0 and Stockfish/Maia PGN protocols.

(6) Conclusion and forward work. Separating unitary core reversibility from irreversible cohomological and Carnot costs, restoring the exact Lambert-W context scaling, and anchoring AI meta-cognition in verifiable behavioral calibration transforms [1–17] into a rigorous, falsifiable computational-physics program.

References

[1] P. De Ceuster, *Quantum Eigenstate Dynamics for Artificial General Intelligence Synthesis*, Zenodo (2025), https://doi.org/10.5281/zenodo.16748240.

[2] P. De Ceuster, *Quantum Eigenstate Dynamics for Artificial General Intelligence Synthesis (Expanded version)*, Zenodo (2025), https://doi.org/10.5281/zenodo.16984175.

[3] P. De Ceuster, *Quantum Sector-Gate Dynamics for Artificial General Intelligence*, Zenodo (2025), https://doi.org/10.5281/zenodo.16968899.

[4] P. De Ceuster, *Interferometric Sector-Gate Quantum Dynamics for AGI: Parity-Driven Tunneling First and the Case for a Topological Core*, Zenodo (2025), https://doi.org/10.5281/zenodo.16983088.

[5] P. De Ceuster, *Interferometric Sector–Gate Nervous Systems: Orchestrating AGI via Parity Fuse Quantum Primitives and Deterministic Java Injection - From Java Injection to Quantum Eigenstates: Engineering an n8n-Style Nervous System for AGI*, Zenodo (2025), https://doi.org/10.5281/zenodo.17204616.

[6] P. De Ceuster, *The Binary–Quantum Bridge for AGI: Information-Preserving Differentiable Programs with a Causal Objective*, Zenodo (2025), https://doi.org/10.5281/zenodo.17490748.

[7] P. De Ceuster, *The Context-Abstraction Limit: Deriving the Critical Token Boundary in Interferometric Sector-Gate AGI*, Zenodo (2025), https://doi.org/10.5281/zenodo.17923954.

[8] P. De Ceuster, *Beyond Artificial Intelligence: The Mathematical Dynamics of General Intelligence - Recursive State-Space Dynamics, Self-Modeling, and Adaptive Goal Formation : From Sector-Gate Dynamics to a Formal Mechanics of AGI*, Zenodo (2026), https://doi.org/10.5281/zenodo.22126076.

[9] P. De Ceuster, *Topological Photonic Acceleration of Artificial General Intelligence: Enhancing Recursive State Dynamics and Parity Compilation via Bulk Gauge Continuity*, Zenodo (2026), https://doi.org/10.5281/zenodo.22165659.

[10] P. De Ceuster, *Intelligence in 2226: Physical Bounds, Neuro-Symbolic Architecture, and the Restructuring of Human Institutions A Realistic Two-Century Perspective*, Zenodo (2026), https://doi.org/10.5281/zenodo.24210205.

[11] P. De Ceuster, *Sub-Landauer Optomagnetic Logic: Non-Thermal Floquet–Chern Switching and Pseudoscalar Photonic Gates for Reversible Neuromorphic Computing*, Zenodo (2026), https://doi.org/10.5281/zenodo.24250267.

[12] P. De Ceuster, *Cohomological and Thermodynamic Guardrails for Causal World Models: A Neuro-Symbolic Architecture for Verifiable Scientific Intelligence*, Zenodo (2026), https://doi.org/10.5281/zenodo.24250689.

[13] P. De Ceuster, *Evaluating Large Language Model Meta-Cognition via the Advanced AGI-AI Self-Awareness Test (AISA-T)*, Zenodo (2025), https://doi.org/10.5281/zenodo.16783096.

[14] P. De Ceuster, *A Follow-up Work to AISA–T: Elucidating the Human Evaluation Rubric for AGI Self-Awareness*, Zenodo (2025), https://doi.org/10.5281/zenodo.16886061.

[15] P. De Ceuster, *Checkmate! Entropy Production, State-Space Dissipation, and Asymptotic Convergence: Revealing Magnus Carlsen's Strategic Chess Trajectories*, Zenodo (2026), https://doi.org/10.5281/zenodo.22037087.

[16] P. De Ceuster, *The Animate and the Inanimate in Pure Mathematics: A Modern Reappraisal of William James Sidis's Underrated Viewpoint*, Zenodo (2025), https://doi.org/10.5281/zenodo.16879071.

[17] P. De Ceuster, *Letter: Pseudoscalar-Mediated Majorana Sector-Gates and Sheaf-Cohomological Core Synthesis for Sub-Landauer Artificial General Intelligence*, Zenodo (2026), https://doi.org/10.5281/zenodo.24291879.

[18] T. Sagawa and M. Ueda, *Generalized Jarzynski equality under nonequilibrium feedback control*, Phys. Rev. Lett. 104, 090602 (2010).

[19] N. Margolus and L. B. Levitin, *The maximum speed of dynamical evolution*, Physica D 120, 188–195 (1998).

[20] S. Bravyi and A. Kitaev, Phys. Rev. A 71, 022316 (2005); D. Litinski, Quantum 3, 205 (2019).

[21] 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).