The Thermodynamic Cost of Answering: a rate–distortion Landauer bound for cyclic intelligence

Daniel Kang

August 2026 — draft v0.3

Abstract

We prove a lower bound on the heat dissipated by a physical machine that answers a stream of queries at a prescribed accuracy. Let $Q \sim D$ be the query, s a bounded score, and $R_D(a) = \inf\{I(Q;\hat{Y}) : \mathbb{E}[s(Q,\hat{Y})] \geq a\}$ the rate–distortion function of the task at accuracy a. For a machine operating cyclically in steady state, *(i)* receiving queries through a read-only, one-way interface and *(ii)* delivering answers by *copy* — the machine still holds a record of what it emitted when the answer departs — the average heat per query satisfies $\langle Q_{\text{diss}} \rangle \geq k_B T \ln 2 \cdot R_D(a)$. Both conditions are necessary, each with an explicit zero-dissipation evasion: if the machine may later read the query register, reversible uncomputation answers for free; if it may hand over its sole physical copy of the answer (move semantics), a swap machine answers for free. Every signal-based channel — wires, radio, light, speech — is a copying channel, so condition (ii) is the physical norm rather than an assumption of convenience. The bound is attained: asymptotically by block-presented batched agents, and within an additive logarithmic term by one-shot agents, via the strong functional representation lemma. A finite-time extension shows attainment requires the zero-power limit: excess dissipation above the frontier is bounded below by a term scaling as the squared correlation displacement over the number of state transitions per query. For temporally correlated query streams the proven floor weakens to $k_B T \ln 2\,[R_D(a) - I(A_0;Q)]_+$, where $I(A_0;Q)$ is the information the machine already holds about the incoming query. The result gives "intelligence per watt" an absolute, task-intrinsic ceiling and a dimensionless efficiency $\eta \leq 1$.

1 Setup

1.1 Task

A *task* is a tuple $(\mathcal{Q}, D, \hat{\mathcal{Y}}, s)$: a finite query alphabet $\mathcal{Q}$ (countable alphabets are discussed in Remark 7), a query distribution D, an answer alphabet $\hat{\mathcal{Y}}$, and a score $s : \mathcal{Q} \times \hat{\mathcal{Y}} \to [0,1]$. If abstention is permitted it is modeled as an explicit symbol of $\hat{\mathcal{Y}}$ with its score defined by s; the accuracy constraint below is always the *unconditional* average over all queries, never an average over answered queries only. All entropies and mutual informations are in bits.

Definition 1 (Intelligence rate)**.** For accuracy level a,

$$R_D(a) \;=\; \inf_{\substack{P_{\hat{Y}|Q}:\\ \mathbb{E}[s(Q,\hat{Y})]\geq a}} I(Q;\hat{Y}). \tag{1}$$

This is the rate–distortion function of D with distortion $d = 1 - s$ at distortion level $1 - a$; it is convex and non-decreasing in a, computable by Blahut–Arimoto for finite alphabets, and

$R_D(a) = 0$ whenever a constant answer achieves accuracy a. Let $a^* = \sup_{P_{\hat{Y}|Q}} \mathbb{E}[s]$, which is attained by a deterministic map when $\mathcal{Q}$ is finite (the objective is linear in the channel). For $a > a^*$ the constraint set is empty, $R_D(a) = +\infty$, and all statements below are vacuously true.

1.2 Physical model

The *agent* A is a classical system with finite or countable state space and finite Shannon entropy at all phase boundaries, coupled to a single *memoryless* heat bath at temperature T, evolving as a Markov jump process whose rates satisfy local detailed balance. Any internal randomness (coins, noise registers) is by definition part of A's state. Heat Q_{diss} is the energy transferred to the bath. Bath Markovianity is load-bearing, not cosmetic: a bath with memory could carry query correlations from the read phase into the emit phase behind the agent's back, and the theorem would fail.

Lemma 1 (Phase inequality)**.** *Let system S evolve on a time interval under Markovian detailed-balance dynamics whose rates may depend on an external random variable Z that is held fixed (read-only) throughout the interval. Then*

$$\beta \langle Q_{\text{diss}} \rangle \;\geq\; \ln 2 \left[H(S|Z)_{\text{start}} - H(S|Z)_{\text{end}} \right], \qquad \beta = 1/k_B T. \tag{2}$$

If the rates do not depend on Z, the inequality also holds with the marginal entropies $H(S)$, even when S is correlated with an inaccessible Z; this marginal form is then the stronger choice, and is the classical content of "erasure without the reference costs the marginal entropy" (del Rio et al., Nature *2011).*

Proof. For each fixed $Z = z$ the process is an ordinary (possibly time-dependent) Markov jump process with detailed-balance rates, for which non-negativity of entropy production gives $\beta\langle Q\rangle_z \geq \ln 2\,[H(S|Z{=}z)_{\text{start}} - H(S|Z{=}z)_{\text{end}}]$ (Esposito and Van den Broeck; Parrondo, Horowitz, Sagawa, *Nat. Phys.* 2015). Averaging over $z \sim P_Z$ yields (2). The frozen read-only Z is the standard infinite-barrier measurement idealization (Sagawa–Ueda; Horowitz–Esposito). When the rates are Z-independent, the marginal law of S evolves autonomously under the master equation and the integral second law applies to it directly. Coupling and decoupling of registers at phase boundaries exchanges work, not heat, and leaves distributions unchanged, so phase-wise composition is legitimate. □

1.3 Cycle protocol

One operating cycle $[t_0, t_3]$ has three phases.

(M) *Read*, $[t_0, t_1]$: a register X prepared in state $Q \sim D$ is presented. The agent's rates may depend on the content of X; the content of X is not altered (read-only). At t_1, X departs.

(E) *Emit*, $[t_1, t_2]$: a blank output register O (fixed initial state o_0) is supplied. The pair (A, O) evolves jointly under bath-coupled dynamics with no dependence on X. At t_2 the answer $\hat{Y} := O(t_2)$ departs with O.

(R) *Reset*, $[t_2, t_3]$: A evolves autonomously (rates depend on no external variable).

Assumptions.

(A1) *Fresh queries*: Q is independent of $A(t_0)$; queries are i.i.d. across cycles, and the query source is not influenced by the agent. Relaxed in Corollary 2.

(A2) *One-way query interface*: X is read-only, departs at t_1, is never re-presented, and its content is the only query-correlated external variable the agent's rates ever see, in any phase.

(A3) *Single memoryless bath, local detailed balance* at temperature T.

(A4) *Steady state*: the marginal distribution of $A(t_3)$ equals that of $A(t_0)$ (per cycle, or amortized over a super-cycle).

(A5) *Copy-semantics emission*: $H(\hat{Y} \mid A(t_2)) = 0$ — when the answer departs, the agent still holds a record determining what it emitted. This is automatic for every signal-based output channel (electrical, optical, radio, acoustic): transmitting modulates a field and leaves the sender's registers intact. It fails only for *move-semantics* delivery, the physical handover of the sole copy (Proposition 2).

(A6) *Accuracy*: $\mathbb{E}[s(Q, \hat{Y})] \geq a$, unconditionally over all queries.

2 Main result

Theorem 1 (Landauer frontier for answering)**.** *Under (A1)–(A6), the average heat dissipated per cycle satisfies*

$$\langle Q_{\text{diss}} \rangle \;\geq\; k_B T \ln 2 \; R_D(a). \tag{3}$$

Proof. Write $H_i(\cdot)$ for entropies at time t_i (all finite by hypothesis). Apply Lemma 1 to each phase with the finest admissible conditioning:

Phase M ($S = A$, $Z = Q$, read-only): $\beta\langle Q_M \rangle \geq \ln 2\, [H_0(A|Q) - H_1(A|Q)]$.

Phase E ($S = (A, O)$, no external dependence; marginal form): $\beta\langle Q_E \rangle \geq \ln 2\, [H_1(A, O) - H_2(A, O)]$.

Phase R ($S = A$, no external dependence; marginal form): $\beta\langle Q_R \rangle \geq \ln 2\, [H_2(A) - H_3(A)]$.

By (A1), $H_0(A|Q) = H_0(A)$. Since $O(t_1)$ is the fixed blank state, $H_1(A, O) = H_1(A)$. By the chain rule and (A5), $H_2(A, O) = H_2(A) + H(\hat{Y}|A(t_2)) = H_2(A)$. By (A4), $H_3(A) = H_0(A)$. Summing:

$$\beta\langle Q_{\text{diss}} \rangle \;\geq\; \ln 2 \left[H_1(A) - H_1(A|Q)\right] \;=\; \ln 2 \; I(A(t_1); Q)\,. \tag{4}$$

During phase E the pair (A, O) evolves under a generator that does not depend on Q, driven by fresh bath noise that is independent of $(Q, A(t_1))$ by (A3), and $O(t_1)$ is deterministic; hence $(A(t_2), \hat{Y})$ is conditionally independent of Q given $A(t_1)$, and $Q \to A(t_1) \to \hat{Y}$ is a Markov chain. (Any private coin the agent might use is part of $A(t_1)$ by the model definition.) The data processing inequality gives $I(A(t_1); Q) \geq I(Q; \hat{Y})$. By (A6) the induced channel $P_{\hat{Y}|Q}$ is feasible for the program (1), so $I(Q; \hat{Y}) \geq R_D(a)$ by definition of the infimum. Combining with (4) proves (3). □

Remark 1 (No converse coding theorem is used)**.** The step $I(Q; \hat{Y}) \geq R_D(a)$ is definitional: $R_D(a)$ is an infimum over feasible channels and the machine's induced channel is feasible. The rate–distortion *coding* theorem enters only on the achievability side.

Remark 2 (Streaming emission)**.** The strict separation of phases M and E is convenience, not load-bearing. If O is present during the read phase and written token by token while X is still visible (autoregressive streaming), applying Lemma 1 to the pair (A, O) with $Z = Q$ in the combined phase gives $\beta\langle Q\rangle \geq \ln 2\,[I(A(t_1); Q) - H(\hat{Y}|A, Q)]$, and (A5) kills the correction. Token-streaming answerers that retain their context obey the same bound.

Remark 3 (Without (A5))**.** Dropping (A5), the same computation yields $\langle Q_{\text{diss}}\rangle \geq k_B T \ln 2\, \Big[R_D(a) - H(\hat{Y} \mid A(t$

This corrected bound is generally *not* tight: a machine that writes $\hat{Y} = m \oplus N$ to O without retaining the noise N pays $k_B T \ln 2\, R_D(a)$ in total even though the corrected bound demands less. Every output-side loophole — move-semantics delivery, oversized registers, unretained randomization — appears in, and only in, this term; the dominant instance is Proposition 2.

3 Both interface conditions are necessary

Proposition 1 (Bennett evasion: necessity of (A2))**.** *Drop (A2) only insofar as the query register X remains* readable *by the agent during phase R. Then for every $a < a^*$ (and $a = a^*$ when the supremum is attained) and every $\varepsilon > 0$ there is an agent meeting the remaining assumptions at accuracy a with $\langle Q_{\text{diss}}\rangle < \varepsilon$.*

Proof sketch. Take deterministic f with $\mathbb{E}[s(Q, f(Q))] \geq a$. The map $(q, \text{blank}) \mapsto (q, f(q))$ is injective, hence implementable by a logically reversible, quasi-statically driven protocol with dissipation $\to 0$. Copy $f(q)$ to O (reversible against the retained scratch; (A5) holds at t_2), deliver, then run the computation backwards against the still-readable X, restoring all internal registers to blank. No erasure without a reference occurs. Note that read access alone suffices; write access to X is not needed. □

Proposition 2 (Swap evasion: necessity of (A5))**.** *Drop (A5) only: allow the emit phase to implement an arbitrary joint permutation of (A, O), including moving the agent's sole record into O. Then for every $a \leq a^*$ attained by some deterministic f and every $\varepsilon > 0$ there is an agent meeting (A1)–(A4), (A6) with $\langle Q_{\text{diss}}\rangle < \varepsilon$.*

Proof sketch. Phase M: with $X = q$ read-only, transport blank internal memory to $m = f(q)$ (deterministic for each q, quasi-static, heat $\to 0$; $H(A|Q)$ remains 0, consistent with Lemma 1). Phase E: *swap* m with the blank O: $(m, o_0) \mapsto (o_0, m)$, a reversible permutation of the joint state space with no X-dependence, heat $\to 0$; O departs carrying m and $A(t_2)$ is blank. Phase R: nothing to erase; $A(t_3) = A(t_0)$ exactly. Total heat $\to 0$ while $R_D(a)$ may be large. □

Remark 4 (What the swap machine means)**.** The swap machine is not a demon trick, and no cost is hidden: handing over the sole copy destroys no correlation, so Landauer charges nothing. The user receives the same physical object O as under copy emission; what changes is that the blank-register supply is consumed as a resource (Bennett; Mandal–Jarzynski), and whoever eventually erases the delivered records pays there. (A5) is thus an accounting statement about *where* the transcript's Landauer cost falls, and a physical statement about the channel: signals copy, couriers swap. A machine that answers by electromagnetic signalling cannot swap-emit, and for it Theorem 1 is unconditional. More generally the bound prices whichever stage of a pipeline creates and retains the Q–$\hat{Y}$ correlation; splitting a system into stages splits the cost, and only the pipeline total is bounded by $k_B T \ln 2\, R_D(a)$.

4 Achievability

Proposition 3 (The frontier is attained)**.** *(i)* One-shot*: for every channel $P_{\hat{Y}|Q}$ achieving accuracy a at rate $I = I(Q;\hat{Y})$ over finite alphabets, there is an agent satisfying (A1)–(A6), with the common randomness W below carried as part of the agent's stationary random state, whose ensemble-average heat obeys*

$$\langle Q_{\text{diss}}\rangle \;\le\; k_B T \ln 2 \left(I + \log_2(I+1) + 4 + \delta\right)$$

for any $\delta > 0$ in the quasi-static limit. (ii) Batched*: agents presented with blocks of n i.i.d. queries in a single read phase attain $\langle Q_{\text{diss}}\rangle/n \to k_B T \ln 2\, R_D(a)$ as $n \to \infty$, taking channels ε-close to the infimum in* (1) *and $\varepsilon \to 0$.*

Proof sketch. (i) By the strong functional representation lemma (Li–El Gamal, *IEEE Trans. IT* 2018), any $P_{\hat{Y}|Q}$ admits $W \perp Q$ and a deterministic g with $\hat{Y} = g(Q,W)$ and $H(\hat{Y}|W) \le I + \log_2(I+1) + 4$. The construction's W is a marked Poisson process; for finite alphabets, truncate to the first N marks, accepting failure probability $\delta_N \to 0$ (an accuracy loss absorbed by δ), so W is a finite register. W is part of A's initial state, never modified, hence stationary and independent of every query; accuracy and heat are ensemble averages over (Q,W). Phase M: compute $m = g(q,W)$ into blank internal registers; for fixed (x,w) this is deterministic transport, logically reversible, cost 0 quasi-statically. Phase E: copy m to O (reversible; (A5) holds since m is retained). Phase R: erase m by a quasi-static protocol whose energy landscape depends on the retained internal reference W; conditional erasure costs $k_B T \ln 2\, H(m|W)$ (matched-landscape protocol: a sudden switch to $E(m,w) = -k_B T \ln p(m|w)$ costs work but no entropy production since the distribution is already Gibbsian, followed by quasi-static compression). (ii) Apply (i) to the product channel $P_{\hat{Y}|Q}^{\otimes n}$, which preserves the per-query unconditional accuracy and has $I(Q^n;\hat{Y}^n) = nI$; the additive overhead is $o(n)$. The block read phase is a real interface concession: the SFRL representation for the product channel is a joint function of the whole block, so all n queries must be simultaneously present; under strictly sequential one-way arrival the agent would have to buffer query copies whose later reference-free erasure destroys the amortization. □

Attainment is quasi-static: any finite answer rate sits strictly above the frontier, quantitatively so by Theorem 2 below. Together, Theorem 1 and Proposition 3 identify $k_B T \ln 2\, R_D(a)$ as the exact thermodynamic price of copy-emitted answering: reversible computing can eliminate every joule above the frontier, and not one below it.

5 The finite-time frontier

Theorem 1 is attained only quasi-statically. This section bounds the cost of speed. Recall the classical speed limit of Shiraishi, Funo, and Saito (*PRL* 121, 070601, 2018): for a Markov jump process on a finite state space over an interval of duration τ, the entropy production Σ (in nats), the time-averaged dynamical activity $\langle\Gamma\rangle$ (expected jumps per unit time), and the total variation displacement obey

$$\Sigma \;\ge\; \frac{\ell^2}{2\,\langle\Gamma\rangle\,\tau} \;=\; \frac{\ell^2}{2\,N}, \qquad \ell := \sum_i \bigl|p_i(\tau) - p_i(0)\bigr|, \tag{5}$$

with $N = \langle\Gamma\rangle\tau$ the expected number of jumps in the interval.

Lemma 2 (Correlation displacement)**.** *Let p be a joint distribution on $\Omega_A \times \Omega_Q$ with mutual information $I(p) \geq R$ bits, and let q be any product distribution on the same space with the same Q-marginal. Then $\ell := \|p - q\|_1$ satisfies*

$$R \;\leq\; 2\,\ell\, \log_2 \frac{|\Omega_A \times \Omega_Q|}{\ell} \qquad \textit{whenever } \ell \leq \tfrac{1}{2}, \tag{6}$$

so, writing $g(\ell) := 2\ell \log_2(|\Omega_{AQ}|/\ell)$, $\ell \geq \ell^*(R) := \inf\{\ell \in (0, 2] : g(\ell) \geq R\}$ unconditionally, with $\ell^*(0) := 0$; for realizable $R > 0$ the infimum is attained.

Proof. Assume $R > 0$ (so $|\Omega_A|, |\Omega_Q| \geq 2$; otherwise the statement is vacuous). $I(q) = 0$. Expanding $I = H_A + H_Q - H_{AQ}$ and using that the Q-marginals coincide, $I(p) \leq |H_A(p) - H_A(q)| + |H_{AQ}(p) - H_{AQ}(q)|$. The L^1 continuity of entropy (Cover–Thomas, Thm. 17.3.3) bounds each term by $\ell' \log_2(|\Omega|/\ell')$ for the respective alphabet and L^1 distance $\ell' \leq \ell$ (marginalization contracts L^1), and $\ell' \mapsto \ell' \log_2(|\Omega|/\ell')$ is increasing on the relevant range ($\ell \leq \frac{1}{2} < 2/e$), giving (6). For $\ell > \frac{1}{2}$: any realizable R obeys $R \leq I(p) \leq \log_2 \min(|\Omega_A|, |\Omega_Q|) \leq \frac{1}{2} \log_2 |\Omega_{AQ}| < \log_2(2|\Omega_{AQ}|) = g(\frac{1}{2})$, so $\frac{1}{2}$ already lies in the constraint set, whence $\ell^*(R) \leq \frac{1}{2} < \ell$ and the conclusion holds; the same bound shows the constraint set is a closed interval bounded away from 0 for $R > 0$, so the infimum is attained. □

Theorem 2 (Finite-time frontier)**.** *Under (A1)–(A6) with finite state spaces, let N_M be the expected number of agent state transitions during the read phase. Then*

$$\langle Q_{\mathrm{diss}} \rangle \;\geq\; k_B T \ln 2\; R_D(a) \;+\; k_B T\, \frac{\ell^*(R_D(a))^2}{2\, N_M}. \tag{7}$$

Proof. Strengthen the phase-M inequality of Theorem 1. With Q frozen, the pair (A, Q) over $[t_0, t_1]$ is itself a Markov jump process on the product space whose jumps are exactly the agent's jumps, with entropy production $\Sigma_M = \ln 2\, \Delta H(A|Q) + \beta \langle Q_M \rangle$ (in nats; the $H(Q)$ term is constant). Apply (5) within each frozen q-sector and use Cauchy–Schwarz over sectors: $\Sigma_M = \mathbb{E}_q[\Sigma_q] \geq \mathbb{E}_q\big[\ell_q^2/(2N_q)\big] \geq \ell^2/(2N_M)$, where ℓ_q is the sector displacement, $\ell = \mathbb{E}_q[\ell_q] = \|p_{AQ}(t_1) - p_{AQ}(t_0)\|_1$ (equality of the average with the joint displacement holds because Q is frozen and $p_{AQ}(t_0)$ is product); the intermediate sector form is sharper and available at no cost. By (A1), $p_{AQ}(t_0)$ is a product distribution with the same Q-marginal as $p_{AQ}(t_1)$, and the proof of Theorem 1 established $I(A(t_1); Q) \geq R_D(a)$; hence Lemma 2 gives $\ell \geq \ell^*(R_D(a))$. Carrying the strengthened phase-M inequality through the telescoping of Theorem 1 adds the excess term verbatim. □

Corollary 1 (Energy–latency–accuracy tradeoff)**.** $\big(\langle Q_{\mathrm{diss}} \rangle - k_B T \ln 2\, R_D(a)\big) \cdot N_M \geq \frac{1}{2} k_B T\, \ell^*(R_D(a))^2$*: excess heat times operations per query is bounded below by a task-intrinsic constant. At fixed activity $\langle \Gamma \rangle$, latency $\tau \geq N_M / \langle \Gamma \rangle$ trades against excess dissipation as $1/\tau$; higher activity (parallelism) relaxes the tradeoff. For $R_D(a) > 0$, $\eta = 1$ is unattainable at any finite operation count within the jump-process model; diffusive or Hamiltonian agents carry no jump count, and for them the optimal-transport route sketched below is the appropriate form.*

Remark 5 (Honest weakness)**.** For macroscopic machines $\ell^*(R) \approx R/(2 \log_2 |\Omega_{AQ}|)$ is astronomically small, so (7) is structural rather than numerically binding: it establishes the *shape* of the frontier — a curve in the (energy, latency) plane approaching $k_B T \ln 2\, R_D(a)$ only as $N_M \to \infty$ — not a practical constraint. Tightening it via optimal-transport entropy-production bounds with graph-distance Wasserstein metrics (Dechant; Van Vu–Saito) in place of total variation is the natural next step, and would replace ℓ^* by a displacement that counts the computation's actual trajectory length.

6 Correlated queries: the value of familiarity

Corollary 2. *Replace (A1) by: queries form a stationary process not influenced by the agent, and the agent may retain arbitrary cross-cycle memory, with (A4) holding in the stationary regime. Then*

$$\langle Q_{\text{diss}} \rangle \;\geq\; k_B T \ln 2 \; [\, R_D(a) \;-\; I(A(t_0); Q)]_+ \,, \tag{8}$$

where D is the marginal query law.

Proof. Identical computation; the only change is $H_0(A|Q) = H_0(A) - I(A(t_0); Q)$. □

Remark 6. This is a lower bound whose slack is unknown: no matching achievability is claimed, and the natural tight object is a conditional rate–distortion function of the query process given the agent's retained side information, which can exceed $R_D(a) - I(A(t_0); Q)$ and remains positive where the displayed floor is vacuous. Qualitatively the direction is meaningful: a machine that already holds predictive information about its user faces a weaker proven floor per query. Note the deployment consequence: in multi-turn interaction the context echoes the machine's previous answers back inside the next query, making $I(A(t_0); Q)$ large; the full $R_D(a)$ floor applies to single-turn, fresh-query operation, and the efficiency anchors below are defined accordingly.

7 The efficiency scale

Definition 2 (Dimensionless intelligence efficiency)**.** For a system with total free-energy input $\langle J \rangle$ per query (all sources: electrical, chemical, and any free energy carried by the interface registers) on task (D, s) at accuracy a,

$$\eta \;=\; \frac{k_B T \ln 2 \; R_D(a)}{\langle J \rangle}.$$

By Theorem 1 and the first law, $\eta \leq 1$ for every one-way, copy-emitting cyclic agent on single-turn tasks. For chemically driven, multi-reservoir systems (brains), (A3) fails as stated, but with all reservoirs isothermal the same proof bounds total entropy production: $T\,\Sigma_{\text{EP}} \geq k_B T \ln 2 \, R_D(a)$ per query, and η against free-energy consumption is the correct form; housekeeping dissipation only widens the gap.

At $T = 300$ K, $k_B T \ln 2 = 2.87 \times 10^{-21}$ J per bit. One concrete anchor, with every intermediate quantity explicit so the factorization multiplies exactly: a laptop-scale 8B-parameter model drafting a 300-token single-turn email at $R_D(a) \sim 10^3$ bits consumes ≈ 100 J and executes $\approx 4.8 \times 10^{12}$ FLOPs; counting one bit-erasure per overwritten register bit (≈ 32 per multiply–accumulate) gives 1.5×10^{14} erasures. Then

$$\frac{1}{\eta} \;=\; \underbrace{\frac{1.5 \times 10^{14}}{10^3}}_{\substack{\text{erasures per task bit} \\ \approx 1.5\times10^{11}\ \text{(algorithmic)}}} \;\times\; \underbrace{\frac{100/1.5 \times 10^{14}}{k_B T \ln 2}}_{\substack{\text{energy per erasure over } k_B T \ln 2 \\ \approx 2.3\times10^{8}\ \text{(hardware, system level)}}} \;\approx\; 3.5 \times 10^{19},$$

consistent with $\eta \approx 3 \times 10^{-20}$. A human brain on the same task (20 W, ~ 5 min, free-energy accounting) has $\eta \approx 5 \times 10^{-22}$. Of the ~ 19.5 orders of magnitude of headroom, roughly 11 are algorithmic (erasures spent per task-relevant bit) and 8.5 are hardware (system-level energy per erasure against the Landauer scale); the larger share of the frontier distance lies in algorithms, though less lopsidedly than device-physics figures of merit (which sit only $10^{3\text{–}4}$ above $k_B T \ln 2$ per switching event) would suggest.

Remark 7 (Regularity and idealizations)**.** Countable alphabets: all statements hold provided the agent's entropy is finite at phase boundaries; Blahut–Arimoto computability and attainment of a^* are finite-alphabet statements, and Proposition 3 then requires the stated truncation. The closed forms used for verification (binary $R = 1 - h(1-a)$; k-ary $R = \log_2 k - h(1-a) - (1-a)\log_2(k-1)$) hold on their standard range $1-a \leq (k-1)/k$. Load-bearing idealizations, each honestly breakable: memoryless bath (A3), frozen read-only registers (infinite-barrier limit), quasi-static attainment (zero throughput), and classical states throughout — with quantum side information, conditional entropies can be negative and the accounting changes (del Rio et al.).

8 Related work

Landauer's principle and its exorcism of Maxwell's demon (Landauer 1961; Bennett 1982) fix the exchange rate between information and heat. Sagawa–Ueda (*NJP* 2013) and Horowitz–Esposito (*PRX* 2014) develop the second law with mutual information for bipartite systems; del Rio et al. (*Nature* 2011) establish that erasure without access to a correlated reference costs the marginal entropy. Still, Sivak, Bell, and Crooks (*PRL* 2012) tie dissipation in driven sensors to non-predictive information; Still (*PRL* 124, 050601, 2020) bounds a cyclic information engine's dissipation by retained non-predictive information, with the information bottleneck emerging as the dissipation-minimizing representation — there a rate–distortion object appears as an *optimizer*, not as a task-indexed universal floor. Theorem 1 differs from both in bounding a task-level, accuracy-constrained quantity with a matching achievability protocol.

Priority. A novelty sweep run after the present derivation was complete surfaced a self-published preprint by A. Vityaz (*Ontology of Transition, Part III*, Zenodo, 21 July 2026, doi:10.5281/zenodo.21473024) which states, for the specialized task of reconstructing an external clock from a reusable physical record, the same converse composition one month earlier: symmetric-Landauer reset work bounded below by the (indirect) rate–distortion function of the reconstruction at prescribed distortion, via Landauer plus data processing, with a no-side-information assumption in the role of our (A2)/(A5), a conditional-side-information extension paralleling Corollary 2, and block-coding attainability. The clock-specialized composition is therefore his. Relative to that work, the present paper treats arbitrary query-answering tasks; derives the cycle bound from phase-wise second-law inequalities rather than axiomatizing the erasure charge; proves both interface conditions *necessary* through explicit zero-dissipation evasions (Propositions 1 and 2), where the earlier work excludes uncomputation and surviving copies by assumption; supplies the one-shot attainability constant via the strong functional representation lemma; adds the finite-time frontier (Theorem 2); and connects the bound to intelligence-per-watt measurement. Niklason (Zenodo, April 2026, doi:10.5281/zenodo.19545467) independently bounds the dissipation of *maintaining* a predictive structure in chemical systems below by its mutual information, in the Still lineage; no accuracy-indexed answering floor appears there. Recent thermodynamic accounts of intelligence propose benchmark-relative bits-per-joule metrics — epiplexity and empowerment per joule (arXiv:2602.05463), which explicitly notes the absence of a universal normalization, and algorithmic-information bounds in the watts-per-intelligence series (arXiv:2504.05328, arXiv:2604.20897). The present bound supplies the universal normalization for the answering setting: $R_D(a)$ is intrinsic to the task, the converse is definitional, and the constant is attained. Quasi-static analog networks with vanishing inference cost (arXiv:2503.09980) realize the reference-present regime of Proposition 1. The intelligence-per-watt measurement program (arXiv:2511.07885) provides the empirical layer this bound turns into a distance from a law of physics.

Verification note

The $R_D(a)$ layer was validated numerically: a Blahut–Arimoto implementation reproduces the closed forms above to 10^{-15}, is invariant under the addition of score-irrelevant answer dimensions (task bits, not text bits), and returns $R = 0$ when a constant answer meets the accuracy level. Theorem 1, both necessity propositions, and the achievability sketch were audited adversarially by three independent referees instructed to refute them; the swap evasion (Proposition 2) was discovered in that audit and promoted from a correction term to a necessity result. Theorem 2 was subsequently audited by a fourth adversarial referee: the speed limit (5) was re-derived independently of the citation, Lemma 2 was tested on 2×10^5 random joint/product pairs and the strengthened phase-M inequality on 60 randomized driven protocols with no violations, and no counterexample to (7) was found; the audit demanded the realizability argument now closing the Lemma's $\ell > \frac{1}{2}$ case and the scope qualifications in the Corollary, and contributed the per-sector strengthening in the proof. Novelty was checked by keyword sweep over arXiv and by citation-graph sweep over OpenAlex (citing sets of Sagawa–Ueda 2013, Still et al. 2012, del Rio et al. 2011, plus phrase searches); the sweep surfaced the Vityaz preprint discussed under Priority above, and nothing else stating a dissipation-versus-rate–distortion floor for cyclic agents.