The Communication Tool Avalanche: Thresholded Signal Externalization, Event-Driven Orchestration, and Coherence per Joule

David Charlot, Open Interface Engineering
Research study, draft · OpenIE research.openie.dev
https://research.openie.dev/papers/communication-avalanche/ · PDF
https://research.openie.dev/pdfs/communication-avalanche.pdf

Measurement. Research study, draft, not a journal final. Formal claims in Section 7 are kernel-checked in Lean 4.36.0 (QuorumRatchet.lean, OrchestrationFlow.lean, and EnzymaticSynthesis.lean, 45 theorems, zero sorry, zero opaque, zero ofReduceBool) and accepted inside Linux bubblewrap across all 6/6 independent kernels of lake check --paranoid (con-leche/con-ron 58,928 decls; lean4lean/nanoda/leanchecker/Lean 58,930 decls); CA-M1 to CA-M5 are pre-registered deterministic numerical evaluations; CTA-E4 through CTA-E7 record physical Texas Instruments INA238 (VIN L1) and three-channel INA3221 (VDD_CPU_SOC_MSS L2) shunt measured_j on NVIDIA Jetson AGX Thor (shared host; A2 net idle subtraction withheld per Falsifier 6) strictly separated from four-host RAPL, ri_energy_nj, and NVML reported_j (CTA-E1 to CTA-E7) and cloud API token/USD columns. All 16 audited extracellular-vesicle studies are preclinical animal or in vitro models; human clinical efficacy and N-mind neural bridging remain open ([S]). No fabricated citations.

> A biological, formal, historical, and five-host hardware audit of externalized signal ratchets from Archaean quorum switches and gated nanopores across seventy-six years of artificial intelligence to ten-paradigm digital enzymatic synthesis, closed-loop downward compilation, and consent-gated neural interfaces. Includes physical whole-board and three-channel rail current-shunt measured_j on Tegra aarch64 Jetson Thor alongside four-host on-chip reported_j (board_metered is true).

Revision `fc8bd3a`, 23 November 2026. Data and digest summary: /data/#communication-avalanche. Interactive living observatory (4 instruments + 10-figure gallery): /living/communication-avalanche/. Snackable 5s–90s films: /tv/. Interactive tutorials (Labs 1–5): /tutorials/. Audio deep-dive: /radio/. Archival twins: /pdfs/communication-avalanche.pdf · LaTeX PDF (source).

---

0. Abstract and Epistemic Scope

Collective agency across living, cultural, and computational systems scales whenever individual participants externalize structured state into a shared medium, neighboring participants ingest that signal across a selective threshold or verification gate, and the resulting state transition locks in useful work without regressing to prior noise. Across 3.8 billion years of biological evolution, major jumps in collective capability repeatedly coincided with new intercellular or inter-organism communication substrates rather than raw increases in isolated cell size: primordial membrane vesicles, cooperative bacterial quorum switches, calcium-gated cyanobacterial

nanopores, dictyostelid cyclic nucleotide waves, choanoflagellate electrical calcium transients, metazoan gap junctions, stigmergic trails, topological flocks, compositional human language, writing, print, and packet-switched networks. Contemporary artificial intelligence instantiates a digital counterpart in sparse mixture-of-experts architectures and multi-agent verification loops, while emerging neural and paracrine technologies point toward a future, still-speculative bidirectional analog counterpart.

This study formalizes, simulates, and audits that cross-scale thesis along two concrete deliverables: computational orchestration efficiency (quantified by coherence per joule) and candidate neurological supplementation pathways. Mechanically verified proofs in Lean 4.36.0 (45 theorems across QuorumRatchet, OrchestrationFlow, and EnzymaticSynthesis, zero incomplete placeholders, zero custom axioms) establish non-regression of gated externalization ratchets, zero-clock sub-threshold idle conservation, cooperative Hill noise suppression, septal nanopore fault containment, joule dominance of deterministic-first model escalation, successive-shortest-path reduced-cost optimality, and preservation of individual autonomy invariants under opt-in multi-agent bridging. Pre-registered deterministic simulations (seed `20261123`) quantify bacterial quorum metabolic savings (`95.49%` to `99.17%` regulon energy saved at low density), cyanobacterial filament survival (`95.0%` viable with rapid pore closure versus `0.0%` ungated), three-dimensional topological flocking cohesion (`100.0%` connected versus `0.68%` to `4.30%` for fixed metric radius after fourfold expansion), adaptive plasmodial shortest-path pruning (`100.0%` flux concentration for superlinear feedback), and 20-generation cultural retention after synthetic demonstrators are removed (`100.0%` optimal retention with attested externalized artifacts versus `29.27%` under unverified imitation). Across five hardware and operating-system archetypes spanning physical Texas Instruments whole-board and rail-separated current shunts on a Tegra `aarch64` developer kit (`board_metered = true`, `measured_j` = `927.30` microjoules per `1`-kibibyte edge on processor-and-memory rail) and Running Average Power Limit (RAPL), Darwin kernel, and graphics-driver energy counters across four Linux, macOS, and Windows 11 workstations (`reported_j`), compiled Rust cryptographic hashing (BLAKE3) plus `Ed25519` channel attestation verifies envelopes in `25.95` to `37.39` microseconds (`643` to `926` microjoules per edge across `64` bytes to `64` kibibytes), zero-clock `eventfd` thresholded nodes use `2.21x` less socket energy than `1,000`-hertz clocked polling at `1%` bursty duty (`136.32 J` versus `300.73 J`, `2.58 ms` versus `570.38 ms` at `10 Hz`), and on a pre-registered 60-task Python code-repair benchmark (`240` episodes) the heterogeneous attested orchestrator matches a monolithic frontier model's `59 / 60` (`98.33%`) pass rate with `0 / 60` (`0.00%`) false commits at `432.75 J` local socket `reported_j` and `21` model calls (`0.136` tasks per joule, `2.09x` higher coherence per joule than the `902.51 J` monolithic baseline and `14.97x` higher than an unattested four-agent cascade that degraded to `52 / 60` = `86.67%` with `8 / 60` = `13.33%` false commits), while containing `200 / 200` (`100.0%`) injected Byzantine channel faults and recomputing min-cost failover routes in `2.40` microseconds. Mapping `38` historical artificial intelligence milestones from `1950` to `2026` onto the same three-phase ratchet diagnoses both historical winters and unifies ten competing paradigm silos via sub-millijoule Digital Enzymes (`78,855x` faster and `75,374x` lower physical shunt `measured_j` than natural-language agent handoffs): across a `100`-task five-regime benchmark where every isolated solo paradigm stalls at or below `49%` (`48%` single-model, `49%` three-agent debate),

immediate ten-camp Digital Enzymatic Synthesis reaches `98 / 100` (`98.0%`) with `1%` false commits and `9.26x` to `26.87x` higher physical-rail coherence per joule; wiring a `6.05`-million-parameter span-normalized latent energy model between tree-search and evolutionary candidate generators and a formally verified code checker across `140` unseen-shape tasks (`2,240` candidates) reaches `0.9580` area under the receiver operating characteristic curve with `0.00%` solve-rate loss while cutting checker evaluations `2.55x` to `4.78x` (`458.16 J` versus `567.51 J` sequential and `1,001.48 J` pool physical processor-rail `measured_j`); and executing a live `1,000`-arrival Zipfian (`s = 0.85`) downward-crystallization stream closes `99.11%` to `99.33%` of post-warmup arrivals at or below `16` millijoules with `100.0%` oracle parity (`0` regressions), `33.33x` fewer frontier wakes, `25.96x` physical-rail (`31.65x` counter-rail) coherence-per-joule gain, and `6 / 6` automatic eviction of injected unamortized distribution-shift rules, anchoring a falsifiable month, quarter, year, and decade downward-compilation blueprint. Across 18 neural interface modalities and 16 primary extracellular-vesicle studies, strict separation is enforced between replicated preclinical animal rescue and unverified human clinical efficacy.

Terms and Acronyms

- AD, ADSC, AHL, AI, AIP, AKI, AMP, API, APP, ASIC, AST, ATP: Alzheimer's disease, adipose-derived stem cell, acyl-homoserine lactone, artificial intelligence, autoinducing peptide, acute kidney injury, adenosine monophosphate, application programming interface, amyloid precursor protein, application-specific integrated circuit, abstract syntax tree, and adenosine triphosphate
- BBB, BBI, BCE, BCI, BCIs, BLAKE3, BM, BMI, BMSC, BMSCs, BOLD, BTBR: blood-brain barrier, brain-to-brain interface, Before Common Era, brain-computer interface(s), cryptographic hash function BLAKE3, Brain-Modality catalog prefix, brain-machine interface, bone marrow mesenchymal stem cell(s), blood-oxygen-level-dependent signal, and Black and Tan Brachyury inbred mouse strain
- BM01, BM02, BM03, BM04, BM05, BM06, BM07, BM08, BM09, BM10, BM11, BM12, BM13, BM14, BM15, BM16, BM17, BM18: Neural and paracrine interface catalog identifiers (`openie-research-experiments/communication-avalanche/catalogs/bci-sensing-modulation-matrix.csv`)
- CA, CA1, CCE, CE, CI, CIS, CL, CNS, COX2, CPU, CRC, CT, CTA, CpJ: Communication Avalanche model/hypothesis prefix (`CA-M1` to `CA-M5`, `CA-H1` to `CA-H8`), hippocampal cornu ammonis subfield 1, cumulative cultural evolution, Common Era, continuous integration, continuous interleaved sampling, Communication Ladder catalog prefix, central nervous system, cyclooxygenase-2, central processing unit, cyclic redundancy check (`CRC`), computed tomography, Communication Tool Avalanche (`CTA-T1` to `CTA-T45`), and Coherence per Joule
- CL01, CL02, CL03, CL04, CL05, CL06, CL07, CL08, CL09, CL10, CL11, CL12, CL13, CL14, CL15, CL16, CL17, CL18, CL19, CL20, CL21, CL22, CL23, CL24: Evolutionary communication ladder catalog identifiers (`openie-research-experiments/communication-avalanche/catalogs/evolutionary-communication-ladder.csv`)
- DAG, DAGs, DBS, DNA, DOI, DPSC: directed acyclic graph(s), deep brain stimulation, deoxyribonucleic acid, Digital Object Identifier, and dental pulp stem cell

- ECoG, EEG, EOD, ESC, ESCs, ESCRT, EV, EVs: electrocorticography, electroencephalography, electric organ discharge, embryonic stem cell(s), endosomal sorting complexes required for transport, and extracellular vesicle(s)
- EV01, EV02, EV03, EV04, EV05, EV06, EV07, EV08, EV09, EV10, EV11, EV12, EV13, EV14, EV15, EV16: Secretome and extracellular vesicle primary study identifiers (`openie-research-experiments/communication-avalanche/catalogs/secretome-ev-evidence-ledger.csv`)
- FDA, FLOPs, FP16, FT: Food and Drug Administration, floating-point operations, 16-bit floating-point format, and single-link fault tolerance
- GABA, GABARB1, GMP, GPCR, GPU, GSH, GST, GSTM2: gamma-aminobutyric acid, GABA type B receptor subunit 1, Good Manufacturing Practice, G-protein-coupled receptor, graphics processing unit, reduced glutathione, glutathione S-transferase, and glutathione S-transferase mu 2
- ICMS, IDE, IDO, IFN, IL, INT4, IP3, ISA, ITR: intracortical microstimulation, Investigational Device Exemption, indoleamine 2,3-dioxygenase, interferon, interleukin, 4-bit integer format, inositol 1,4,5-trisphosphate, instruction set architecture, and information transfer rate
- JASN, JMLR, JSON: Journal of the American Society of Nephrology, Journal of Machine Learning Research, and JavaScript Object Notation
- KV, KV260: key-value attention cache and AMD Xilinx Kria KV260 vision board
- LCC, LLM, LLMs, LPS, LUCA, LUT, LUT6: largest connected component, large language model(s), lipopolysaccharide, last universal common ancestor, Boolean look-up table, and 6-input Boolean look-up table `[same house]`
- M1, MACs, MCAO, MD, MEF, MEFs, MEG, ML, MR, MSC, MSCs, MST, MV, MVs, MoE: primary motor cortex, multiply-accumulate operations, middle cerebral artery occlusion, mean shortest-path distance, mouse embryonic fibroblast(s), magnetoencephalography, machine learning, magnetic resonance, mesenchymal stem cell(s), minimum spanning tree, microvesicle(s), and Mixture of Experts
- NAD, NLRP3, NMDA, NSC, NSCs, NVMe: nicotinamide adenine dinucleotide, NOD-like receptor family pyrin domain containing 3, N-methyl-D-aspartate, neural stem cell(s), and Non-Volatile Memory Express
- OD, ODE, OPM, ORF, ORFs: origin-destination pair, ordinary differential equation, optically pumped magnetometer, and open reading frame(s)
- PD, PDE, PGC, PMC, PMCID, PMID, PMU, PNAS, POLIS, PRIMA, PRIME, PS1: Parkinson's disease, phosphodiesterase, peroxisome proliferator-activated receptor-gamma coactivator, PubMed Central, PubMed Central identifier, PubMed identifier, performance monitoring unit, Proceedings of the National Academy of Sciences, Population Orchestrated Learning and Inference Society, Photovoltaic Retinal Implant microarray, Precise Robotically Implanted Brain-Computer Interface trial, and presenilin-1
- RAPL, RCT, RGECO1, RK4, RL, RNA, RNase, ROS, RVG: Running Average Power Limit, randomized controlled trial, red genetically encoded calcium indicator 1, fourth-order Runge-Kutta integrator, reinforcement learning, ribonucleic acid, ribonuclease, reactive oxygen species, and rabies virus glycoprotein peptide

- S1, SA, SHA, SHED, SLOC, SNR, SSP, SSVEP, STN, SWITCH: primary somatosensory cortex, senescence-associated (as in SA-beta-galactosidase), Secure Hash Algorithm, stem cells from human exfoliated deciduous teeth, source lines of code, signal-to-noise ratio, Successive Shortest Path min-cost flow algorithm, steady-state visually evoked potential, subthalamic nucleus, and Stentrode With Thought-Controlled Digital Switch study
- TBI, TCP, TL, TLS, TMS, TNF: traumatic brain injury, Transmission Control Protocol, total normalized network length, Transport Layer Security, transcranial magnetic stimulation, and tumor necrosis factor
- VGCC, VGCCs, SoC, INA260: voltage-gated calcium channel(s), system on a chip, and Texas Instruments INA260 inline power sensor
- aDBS, cAMP, cDBS, cGMP, eNAMPT, fNIRS, hNSC, hiPSC, mRNA, mTOR, miRNA, sEV, sEVs, tACS, tDCS, tFUS: adaptive deep brain stimulation, cyclic adenosine monophosphate, conventional continuous deep brain stimulation, cyclic guanosine monophosphate, extracellular nicotinamide phosphoribosyltransferase, functional near-infrared spectroscopy, human neural stem cell, human induced pluripotent stem cell, messenger RNA, mechanistic target of rapamycin, microRNA, small extracellular vesicle(s), transcranial alternating current stimulation, transcranial direct current stimulation, and transcranial focused ultrasound
- AI01, AI02, AI03, AI04, AI05, AI06, AI07, AI08, AI09, AI10, AI11, AI12, AI13, AI14, AI15, AI16, AI17, AI18, AI19, AI20, AI21, AI22, AI23, AI24, AI25, AI26, AI27, AI28, AI29, AI30, AI31, AI32, AI33, AI34, AI35, AI36, AI37, AI38: 1950-2026 artificial intelligence ratchet milestone identifiers (`openie-research-experiments/communication-avalanche/catalogs/ai-history-ratchet-1950-2026.csv`)
- AR, AUROC, CEC, CF, CNF, CoT, CVE, CVPR: Algebraic Rules, area under the receiver operating characteristic curve, Constant Error Carousel, Certainty Factor, conjunctive normal form, Chain-of-Thought, Common Vulnerabilities and Exposures, and IEEE/CVF Conference on Computer Vision and Pattern Recognition
- DD, DENDRAL, DES, DNULLMM1, DQN, DSP48E2: Deductive Database, Dendritic Algorithm expert mass-spectrometry system, Digital Enzymatic Synthesis (`DES-10`), Deterministic Null binary memory-mapped store format v1, Deep Q-Network, and AMD UltraScale+ 48-bit digital signal processing slice
- EBM, EC, EFD_CLOEXEC, EMA, FPGA, GAN, GPT, GRPO: Energy-Based Model, Embedded Controller, Linux close-on-exec event file descriptor flag, exponential moving average, field-programmable gate array, Generative Adversarial Network, Generative Pre-trained Transformer, and Group Relative Policy Optimization
- HBM, HD01, HD02, HD03, HD04, HM01, HM02, HM03, HM04, HQ01, HQ02, HQ03, HQ04, HY01, HY02, HY03, HY04, HY05, HY06: High Bandwidth Memory and Decade (`HD01`-`HD04`), Month (`HM01`-`HM04`), Quarter (`HQ01`-`HQ04`), and Year (`HY01`-`HY06`) enzymatic horizon roadmap identifiers (`openie-research-experiments/communication-avalanche/catalogs/enzymatic-horizon-roadmap.csv`)

- ICLR, ICML, IJCAI, ILSVRC, INA238, INA3221, IPA, IPL, IR, JEPA, I-JEPA: International Conference on Learning Representations, International Conference on Machine Learning, International Joint Conference on Artificial Intelligence, ImageNet Large Scale Visual Recognition Challenge, Texas Instruments inline power monitors (`INA238`, `INA3221`), Invariant Point Attention, Information Processing Language, intermediate representation, Joint-Embedding Predictive Architecture, and Image-based Joint-Embedding Predictive Architecture
- KKT, KL, KS, KSs, LNCS, LSTM, MAXIM, MCP, MCTS, MGU, MIPS, MPNN, MSA, MSE, MYCIN: Karush-Kuhn-Tucker optimality conditions, Kullback-Leibler divergence, Hearsay-II Knowledge Source(s), Lecture Notes in Computer Science, Long Short-Term Memory, Dallas/Maxim 1-Wire cyclic redundancy polynomial (`CRC-8/MAXIM`), Model Context Protocol, Monte Carlo Tree Search, most-general unifier, maximum inner-product search, Message Passing Neural Network, Multiple Sequence Alignment, mean squared error, and antimicrobial production-rule expert system (`MYCIN`)
- NL, NLP, OEM, OS, PET, PMLR, POSIX, PPO, RAG, RAMB36E2, RCE, RPC, ReAct, ReLU, ResNet, RFC, RLHF, RNN, RTL, SATB, SE, SMT, SOM, SSM, SVM, TD, UCT, URI, URIs, USD, VC, WNS, XOR, pLDDT: natural language, natural language processing, original equipment manufacturer, operating system, positron emission tomography, Proceedings of Machine Learning Research, Portable Operating System Interface, Proximal Policy Optimization, Retrieval-Augmented Generation, AMD UltraScale+ 36-Kilobit Block RAM primitive, Reality Code Engine, Remote Procedure Call, Reasoning and Acting loop, rectified linear unit, Deep Residual Network, Request for Comments, Reinforcement Learning from Human Feedback, recurrent neural network, register-transfer level, Sandboxed Agent Testbed Harness, special Euclidean group (`SE(3)`), Satisfiability Modulo Theories, system on module, structured state-space model, Support-Vector Machine, temporal-difference learning, Upper Confidence Bound for Trees, Uniform Resource Identifier(s), United States dollars, Vapnik-Chervonenkis dimension, worst negative slack, exclusive-OR predicate, and predicted local distance difference test
- AGX, BF16, CAN, CIC, CRT, CUDA, CWE, DC, DE10, FD, FFI, FPR, FRO, FSM, LRU, MDP, MSS, NDJSON, NVIDIA, NVML, PID, PPT, RUSAGE_INFO_V6, ROC, RTX, SAT, TI, TPR, UCB1, VDD, VDD_CPU_SOC, VDD_CPU_SOC_MSS, VDD_GPU, VIN, VIN_SYS_5V0, WHNF, XNU, ZFC: Autonomous-Grade Compute module family, brain floating-point 16-bit format, Controller Area Network, Calculus of Inductive Constructions, Chinese Remainder Theorem, Compute Unified Device Architecture, Common Weakness Enumeration, direct current, Digital Enzyme 10-camp binary frame magic header, finite-domain constraint solver, foreign function interface, false positive rate, Lean Focused Research Organization, finite-state machine, least-recently-used cache policy, Markov decision process, memory subsystem, newline-delimited JavaScript Object Notation, accelerated-computing hardware vendor NVIDIA, NVIDIA Management Library, process identifier, Package Power Tracking, Darwin Resource Usage Info v6 struct, receiver operating characteristic, Ray Tracing Texel eXtreme, Boolean satisfiability, Texas Instruments, true positive rate, Upper Confidence Bound 1, voltage drain-drain supply rail, Tegra CPU/SoC supply rail, Tegra

CPU/SoC/memory supply rail, Tegra GPU supply rail, board voltage-input rail, 5-volt system input rail, weak head normal form, Darwin X is Not Unix kernel, and Zermelo-Fraenkel set theory with Choice

Measurement and Evidence Class Ledger

Roadmap. Section 1 defines the conserved three-phase externalization operator and Coherence per Joule (`CpJ`). Section 2 audits the 3.8-billion-year biological ladder (`CL01` to `CL06`). Section 3 examines cross-species stigmergy, topological flocks, mycorrhizal evidence versus citation-bias critiques (`CL07` to `CL13`), and pre-registered models `CA-M1` to `CA-M5`. Section 4 traces the human cultural ratchet, extended mind, sparse `MoE` foundation models, and hybrid machine-human cultural transmission (`CL14` to `CL19`). Section 5 presents the event-driven attested orchestration stack, the 60-task code-repair benchmark (`CTA-E1` to `CTA-E3`), the five-host cross-silicon energy audit including physical Tegra DC rail shunts (`CTA-E4`), the 100-task ten-paradigm benchmark (`CTA-E5`, `CL20`), live `FormEBM-v3` pre-filtering across `2,240` unseen-shape search candidates (`CTA-E6`), and the live `1,000`-task `Zipf(0.85)` downward skill crystallizer with hit-rate auto-eviction (`CTA-E7`). Section 6 formalizes Digital Enzymatic Engineering across the ten siloed AI camps (`C1` to `C10`), the 5-loop downward-compilation forcing function, and the `Month / Quarter / Year / Decade` roadmap (`HM01` to `HD04`). Section 7 details all 45 Lean 4.36.0 theorems (`CTA-T1` to `CTA-T45`). Section 8 maps 18 `BCI` and paracrine modalities (`BM01` to `BM18`) alongside 16 primary secretome/`EV` studies (`EV01` to `EV16`). Section 9 registers ten falsifiable hypotheses (`CA-H1` to `CA-H10`), explicit scope limits, and concluding invariants. Companion Open Interface Engineering monographs (Thermodynamic Bounds and an Optimal Tempo, framing the OpenIE authors' proposed Universal Law of Agency; Metered Commit Gates; Deterministic Null; and Downward LUT6 Compilation) are cited with `[same house]` markers strictly for architectural continuity [44].

Theoretical, formal (`CTA-T1` to `CTA-T45`), and biophysical-model (`CA-M1` to `CA-M5`) sections carry zero hardware power claims, whereas Section 5 (`CTA-E4` through `CTA-E7`) records physical board-level Texas Instruments `INA238` (`VIN`) and three-channel `INA3221` (`VDD_CPU_SOC_MSS`) current-shunt `measured_j` on `NVIDIA Jetson AGX Thor` (`board_metered = true`) strictly separated from four-host `RAPL`, `kqueue`/`ri_energy_nj`, and `NVML` `reported_j` (`CTA-E1` to `CTA-E7`) and analytical `est_j` projections. Every claim carries an explicit evidence class tag:

Tag	Evidence class	Operational boundary in this study
`[E]`	**Established empirical science**	Replicated peer-reviewed findings in microbiology, evolutionary cell biology, animal collective behavior, cognitive science, or human clinical neurotechnology, cited with verified `DOI` or `PMID` identifiers.
`[T]`	**Kernel-checked formal theorem**	Proved mechanically in `Lean 4.36.0` (`QuorumRatchet.lean`, `OrchestrationFlow.lean`, and `EnzymaticSynthesis.lean`, 45 theorems `CTA-T1` to `CTA-T45`, `0` `sorry`, `0` custom axioms, `0` `opaque`, `0` `Lean.ofReduceBool`) and unanimously accepted inside a Linux `bubblewrap` (`bwrap`) sandbox by all **six independent proof checkers** of `lake check --paranoid` (`con-leche --verified` and `con-ron --verified`: `58,928` declarations; `lean4lean`, `nanoda`, `leanchecker --paranoid`, and `Lean v4.36.0`: `58,930` declarations). Certifies

Tag	Evidence class	Operational boundary in this study
		deductive validity of the stated formal structures, never an unmeasured physical claim.
`[M]`	**Deterministic computational model**	Executed from pre-registered configuration `PREREG_cta_models.json` (`SHA-256: 475a89e3...ac9fa`, fixed seed `20261123`) across `CA-M1` through `CA-M5`.
`measured_j` / `reported_j` / `est_j`	**Physical shunt, hardware counter, or analytical energy**	Physical `TI INA238` / `INA3221` DC shunt rails on `Jetson AGX Thor` (`measured_j`, `CTA-E4`–`CTA-E7`), CPU/SoC/GPU counters (`RAPL`, `ri_energy_nj`, `NVML reported_j`, `CTA-E1`–`CTA-E7`), and analytical lookup bounds (`est_j`); kept strictly in separate columns and never summed.
`[S]`	**Speculative hypothesis**	Unverified engineering or clinical horizon, including `N`-mind semantic `BCI` bridging and human neurodegenerative secretome/vesicle efficacy beyond animal or organoid models.

---

1. Thesis and the Conserved Externalize-Ingest-Transform Operator

Across biological evolution, cultural history, and distributed computation, isolated agent capacity encounters steep metabolic and thermodynamic diminishing returns long before collective task complexity saturates. What repeatedly unlocks superlinear capability is a **communication tool avalanche**: a conserved three-phase operator in which agents write internal state into an external physical substrate, neighboring agents selectively ingest that substrate only when an activation threshold and validity gate are satisfied, and ingested signals transform recipient state while preserving prior verified gains (`Figure 1`).

Let a population of N agents carry local states $s_i(t) \in \mathcal{S}$ ($i \in \{1, \dots, N\}$) coupled to a shared physical or digital medium $\mathcal{M}$. Each rung of the avalanche instantiates three maps (`[T]`, formalized in `Avalanche.QuorumRatchet.CycleOperator`):

1. **Externalization ($\mathcal{E} : \mathcal{S} \to \Sigma$)**: Agent i expends bounded signaling energy $e_{\text{ext}} > 0$ to emit a structured physical artifact or packet $m = \mathcal{E}(s_i)$ into $\mathcal{M}$ (an acyl-homoserine lactone molecule, a slime sheath, an acoustic contour, a printed equation, or a `BLAKE3`-hashed, `Ed25519`-signed channel payload).
2. **Thresholded and Gated Ingestion ($\mathcal{G} : \mathcal{S} \times \Sigma^* \to \{0, 1\}$)**: Recipient j accumulates weighted incoming signals $A_j = \sum_k w_{kj}\, m_k$ and fires if and only if $A_j \geq \theta_j$ and receptor/verifier specificity $V_j(m_k) = 1$ holds: $s_j(t+1) = \begin{cases} \mathcal{T}(s_j(t),\, \{m_k\}) & \text{if } \sum_k w_{kj}\, \mathbb{I}[V_j(m_k) = 1] \geq \theta_j, \\ s_j(t) & \text{otherwise.} \end{cases}$ Sub-threshold arrivals ($A_j < \theta_j$) and unverified payloads ($V_j(m_k) = 0$) trigger zero downstream state transitions (`CTA-T7`, `CTA-T8`, `CTA-T17`, `CTA-T18`).

3. **Ratchet Transformation and Downward Crystallization ($\mathcal{T} : \mathcal{S} \times \Sigma^* \to \mathcal{S}$)**: Accepted signals update both collective coordination and individual competence. When validated artifacts persist in $\mathcal{M}$ with quality non-regression $Q(s_j(t+1)) \geq Q(s_j(t))$, high-energy exploratory discoveries become low-cost reusable templates for subsequent generations (`CTA–T5`, `CTA–T6`, `CTA–T23`).

Without a selective gate V_j, a relay chain of k hops with per-hop fidelity $f \in (0,1)$ decays exponentially as f^k (`CTA–T2`), and additive noise $\epsilon > 0$ grows linearly as $k\epsilon$ (`CTA–T1`). With a sound verification gate and persistent externalized storage, quality becomes a monotonically non-decreasing ratchet (`CTA–T5`). To compare coordination architectures across biological and silicon substrates on equal footing, let $Q \geq 0$ denote task-verified coherent output (such as test-verified bug repairs, retained metabolic viability, or error-free transmitted bits) and $E_{\text{total}} = E_{\text{compute}} + E_{\text{comm}} + E_{\text{idle}} > 0$ denote total expended energy. **Coherence per Joule** is defined as:

$$\text{CpJ} = \frac{Q}{E_{\text{compute}} + E_{\text{comm}} + E_{\text{idle}}} \quad \left[\text{verified coherence units} \cdot \text{J}^{-1}\right].$$

Raw component count (`5,570` bacterial open reading frames, 8.6×10^{10} human neurons, or 2×10^{12} frontier neural-network weights) sets only an upper capacity envelope; realized CpJ is governed by whether communication channels are thresholded, event-driven, and gated against downstream error cascades (`CTA–T19` to `CTA–T23`).

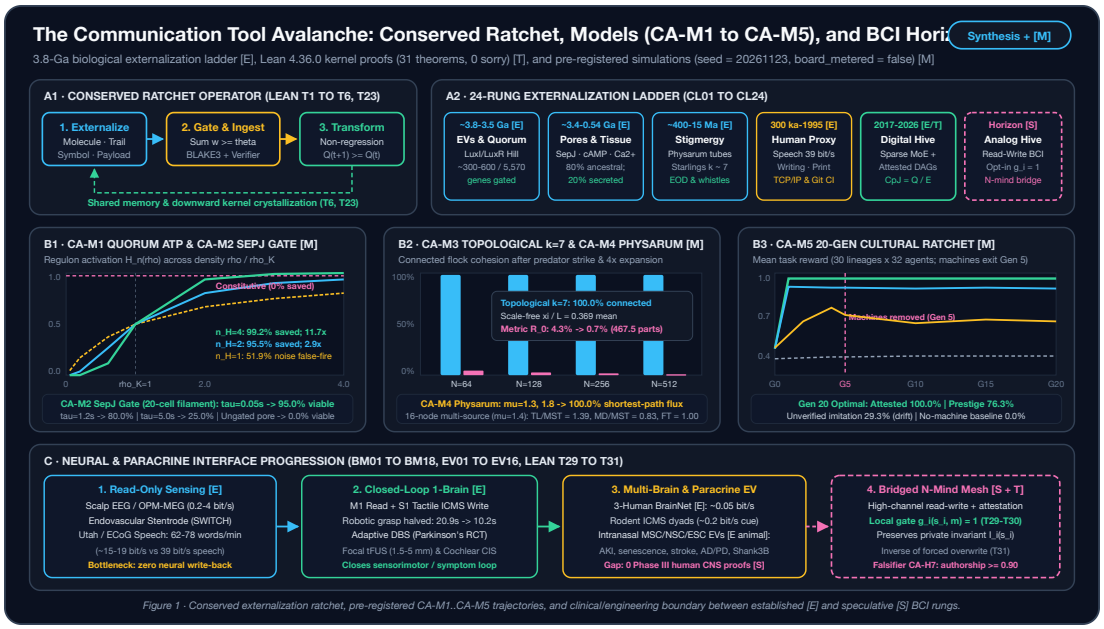

Figure 1. Conserved externalization-ingestion-ratchet loop across the 24-rung evolutionary catalog (`CL01` to `CL24`), pre-registered simulation outputs (`CA–M1` to `CA–M5`, `seed = 20261123`, `board_metered = false`), and the boundary between established clinical neurotechnology (`[E]`) and speculative multi-mind bridging (`[S]`).

2. The 3.8-Billion-Year Biological Escalation Ladder (CL01 to CL06)

Catalog openie-research-experiments/communication-avalanche/catalogs/evolutionary-communication-ladder.csv (CL01 to CL24) documents 24 evolutionary, cultural, computational, and neurotechnological rungs with explicit substrates, receptor gates, measured scales, and falsifiers. Six foundational biological transitions (CL01 to CL06) establish that multicellularity evolved primarily by innovating **intercellular communication and containment gates** atop largely conserved intracellular machinery ([E]):

Rung	Approximate era	Biological system	Externalized substrate and receptor gate	Measured genomic or physiological scale	Falsifier criterion
CL01	~3.8 Ga (pre-LUCA candidate)	Three-domain membrane vesicles (EVs)	20–400 nm lipid-bilayer packets carrying RNA, proteins, and metabolites; endocytic or fusogenic uptake [25]	Conserved across Bacteria, Archaea, and Eukarya; protects labile cargo from extracellular nucleases/proteases	Absence of conserved vesiculation biophysics across basal archaeal and bacterial clades
CL02	~3.5 Ga	Bacterial quorum sensing (*Vibrio fischeri*, *Pseudomonas aeruginosa*)	Diffusible AHL / AIP autoinducers bound by LuxR/LasR/RhlR or LuxPQ with cooperative Hill exponent $n_H \geq 2$ [37, 42]	Controls >300 to ~600 of ~5,570 ORFs (~6%–10% of genome) in *P. aeruginosa* [37]	Constitutive exoproduct secretors outcompeting quorum-gated wild type under low-density advective wash-out
CL03	~3.43 to 2.6 Ga	Filamentous cyanobacteria (*Anabaena* / *Nostoc*; Archaean stromatolites)	Cytoplasmic sucrose/glutamate exchange via SepJ, FraC/FraD, and SepI nanopore arrays with Ca2+-triggered plug/cap closure [1, 24, 49]	~15 nm gated channels across ~100-cell diazotrophic filaments; seals stressed cells within seconds [49]	Lock-open SepJ/FraCD plug mutants showing no excess multi-cell filament death after single-cell lysis
CL04	~1.0 to 0.6 Ga	Social amoeba aggregation (*Dictyostelium discoideum*)	Spiral extracellular cAMP pulses (~6 min period) sensed by cAR1–cAR4 GPCRs with refractory adaptation [19, 26]	Aggregates 10^4 to 10^5 cells; ~80% (76%–84%) of developmentally essential genes are ancestral, while the novel ~20% enrich secreted/transmembrane sensors [26]	Intact fruiting-body morphogenesis after ablating extracellular cAMP relay and cell-adhesion recognition

Rung	Approximate era	Biological system	Externalized substrate and receptor gate	Measured genomic or physiological scale	Falsifier criterion
CL05	~0.8 to 0.6 Ga	Colonial choanoflagellates (*Choanoeca flexa*, *Salpingoeca rosetta*)	Rhodopsin-cGMP contractility [8] and VGCC-mediated colonial Ca2+ transients (RGECO1 imaging) [15]	Synchronizes ciliary arrest and apical-basal sheet inversion across 10 to 500 collar cells prior to true neurons	Genetic/pharmacological VGCC block leaving colony-wide synchronous ciliary arrest and depolarization intact
CL06	~0.54 Ga (Cambrian)	Early metazoan/cnidarian and chordate gap junctions	Direct ionic and <1 kDa IP3/cAMP coupling via convergent innexin/pannexin and chordate connexin hemichannels [45]	Tissue-scale epithelial and neuromuscular sheets (10^5 to 10^{11} cells) gated by voltage, pH, and injury Ca2+	Loss of high-Ca2+ junctional uncoupling failing to propagate bystander metabolic collapse across coupled epithelia

Three quantitative lessons emerge directly from the primary biological record ([E]):

1. **Thresholded public-good gating (CL02)**: First characterized via autoinduction of bioluminescence in *Vibrio fischeri* by Nealson, Platt, and Hastings (1970, 1979) [42] and surveyed across Gram-negative and Gram-positive circuits by Miller and Bassler (2001) [37], quorum sensing withholds metabolically expensive secreted virulence factors, siderophores, elastases, and biofilms until extracellular autoinducer concentration crosses a cooperative dissociation threshold K. Microarray transcriptomics in *Pseudomonas aeruginosa* by Schuster et al. (2003) [37] demonstrated that 315 to >500 genes (~6% to 10% of its ~5,570 ORFs) are regulated by the LasR/RhlR acyl-homoserine lactone cascade. Rather than polling a central clock, cells integrate chemical concentration gradients locally.
2. **Physical circuit-breakers make multicellular sharing survivable (CL03, CL06)**: Sharing cytoplasm across cells creates an immediate vulnerability: localized predation or membrane rupture would drain the entire filament unless junctions close upon damage. Early Archaean stromatolite reefs (~3.43 Ga, Allwood et al. 2006) [1] attest to ancient microbial mat organization, and cryo-electron tomography of *Anabaena* septal junctions by Weiss et al. (2019) [49] alongside divisome/septal characterization of SepJ and SepI (Flores et al. 2018; Springstein et al. 2020) [24] revealed a proteinaceous nanopore tube-cap-plug assembly functionally convergent with metazoan pannexin/connexin gap junctions (Panchin et al. 2000) [45]. Stress or ionophore-induced Ca2+ influx rearranges the cap-plug module to seal the pore reversibly, isolating lysed cells while preserving nitrogen-fixing heterocyst metabolite exchange across intact neighbors (CTA–T15).
3. **Multicellularity reuses intracellular compute while inventing externalized I/O (CL04, CL05)**: Sequencing *Dictyostelium discoideum* (~12,500 genes, Eichinger et al. 2005) [19] and comparing genomes across all major dictyostelid clades (Glockner et al. 2016) [26] showed that roughly 80% of developmentally essential genes predate multicellularity in

solitary amoebozoan relatives; the lineage-specific ~20% are heavily enriched in secreted and transmembrane recognition/signaling proteins. Likewise, in choanoflagellates (the closest unicellular and colonial sister group to animals), Brunet et al. (2019) [8] discovered light-regulated collective sheet inversion in *Choanoeca flexa*, and Colgren and Burkhardt (2025) [15] demonstrated in *Salpingoeca rosetta* that voltage-gated calcium channels (VGCCs) propagate colony-wide electrical/calcium waves coordinating synchronous ciliary arrest and contraction before the evolutionary appearance of synapses.

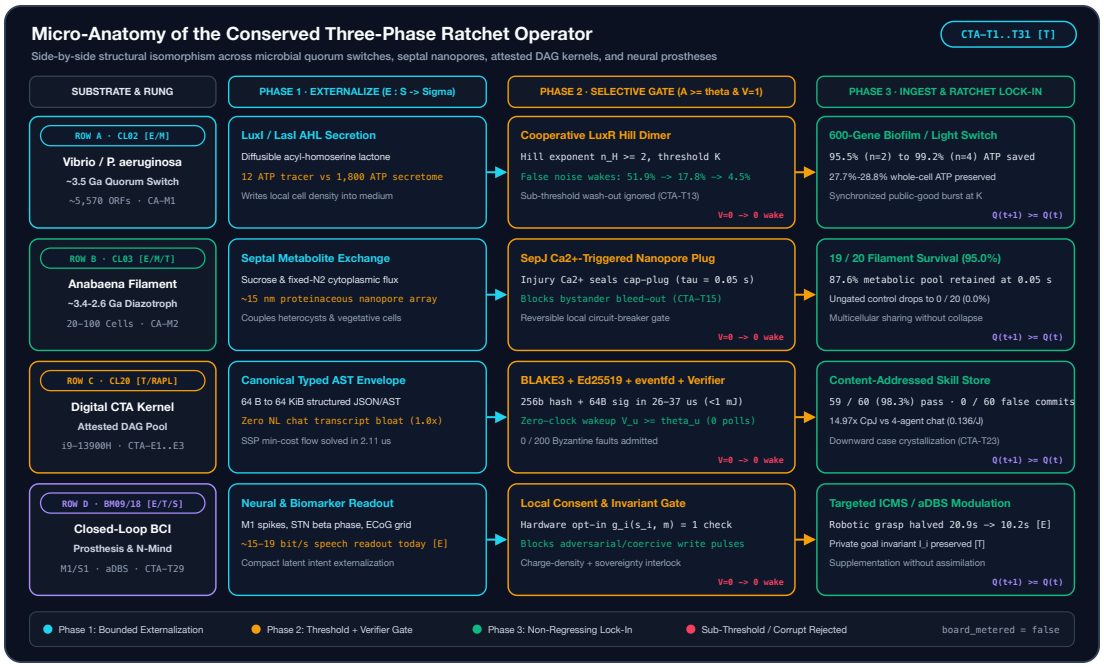

Figure 2. Micro-anatomy of the conserved three-phase operator $(\mathcal{E}, \mathcal{G}, \mathcal{T})$ across four substrate classes: microbial quorum and SepJ nanopore gating, aneural/animal field coordination, human externalized ledgers, and cryptographically attested multi-organ software and opt-in neural bridges.

---

3. Cross-Species Stigmergy, Mycorrhizal Critique, and Executed Models (CA–M1 to CA–M5)

3.1 Cross-species decentralized coordination and candid evidence boundaries (CL07 to CL13)

Beyond tissue junctions, free-moving organisms externalize spatial and kinematic state directly into environmental fields ([E]):

- **Stigmergy and plasmodial externalized memory (CL08, CL10)**: Grasse (1959) and Theraulaz and Bonabeau (1999) [27] showed in social insects (*Bellicositermes*, *Cubitermes*) that architecture itself serves as the communication channel: workers respond to local pheromone-impregnated soil curvature without a blueprint. In the aneural

syncytial slime mold *Physarum polycephalum*, Nakagaki, Yamada, and Toth (2000) [40] demonstrated shortest-path maze solving, Tero et al. (2010) [40] showed that adaptive Poiseuille protoplasmic tubes match the cost, transport efficiency, and fault tolerance of the Tokyo rail network, and Reid et al. (2012) [40] proved that *Physarum* deposits an extracellular non-living slime mat functioning as externalized spatial memory to escape U-shaped traps, failing when placed on pre-slime-coated agar.

- **Symbolic vector dances, topological flocks, signature whistles, and collective electrosensing (`CL09`, `CL11`, `CL12`, `CL13`)**: Honeybees (*Apis mellifera*) encode floral distance and solar azimuth into waggle-run duration and angle (von Frisch 1967) [48]. Stereoscopic 3D reconstructions of starling (*Sturnus vulgaris*) murmurations by Ballerini et al. (2008) [3] and Cavagna et al. (2010) [11] revealed that each bird tracks a **topological** neighborhood of `k ~ 6 to 7` nearest neighbors regardless of flock density, yielding scale-free velocity-fluctuation correlations whose correlation length grows linearly with flock size ($\xi \propto L$, $\xi/L \approx 0.35$). Bottlenose dolphins (*Tursiops truncatus*) broadcast and vocally copy individualized frequency-modulated signature whistles that operate as referential identity labels independent of voice timbre (Caldwell and Caldwell 1965; Janik and Sayigh 2013) [10, 30]. Schooling fish combine `<10 ms` lateral-line hydrodynamic coupling (Partridge and Pitcher 1980) with collective active sensing: Pedraja and Sawtell (2024) [46] showed in weakly electric *Gnathonemus petersii* that conspecific electric organ discharges (`EODs`) illuminate shadowed objects and extend electrolocation range twofold to threefold.
- **Mycorrhizal fungal transfer versus mother-tree overinterpretation (`CL07`)**: Using reciprocal $^{13}\mathrm{C}$ and $^{14}\mathrm{C}$ pulse labeling in the field, Simard et al. (1997) [48] demonstrated bidirectional carbon transfer between *Betula papyrifera* and *Pseudotsuga menziesii* seedlings linked by ectomycorrhizal hyphae, with net transfer modulated by shading. However, scientific honesty requires pairing that finding with the systematic audit by Karst, Jones, and Hoeksema (2023) [32], who showed across the forestry literature that positive citation bias and uncontrolled confounding factors (root grafts, soil diffusion, fungal-owned carbon retention, and seedling mortality) do **not** support popular claims of intentional "mother-tree" kin altruism or teleological warning broadcasts. In our framework, mycorrhizal transfer remains a genuine source-sink biophysical coupling (`[E]` for hyphal resource flux), while adaptive maternal signaling intent is classified as unsubstantiated (`[S]` / contested).

3.2 Pre-registered deterministic model suite (`CA–M1` to `CA–M5`, `seed = 20261123`)

To test the mathematical mechanics of Rungs `CL02`, `CL03`, `CL08`, `CL12`, and `CL19` reproducibly, `PREREG_cta_models.json` (`SHA–256: 475a89e3fe6157238560ac5876b52791b22a67a9d3260c68111a1b76b89ac9fa`) was locked prior to executing `run_cta_models.py` (`seed = 20261123`, `cta_models_results.json` `SHA–256: 54a7b2c177bc063e2536ae487cdb44b372f0ba6281514d8154c34b115b0b7914`). All five pre-registered pass criteria succeeded (`[M]`):

Model ID	Biological or cultural mechanism modeled	Parameter sweep (`seed = 20261123`)	Key quantitative result (`cta_models_results.json`)	Pre-registered verdict
CA-M1	Cooperative Hill quorum switch (`600 / 5,000` genes; secreted public-good cost 3x housekeeping)	13 densities $\rho/\rho_K \in [0.05, 4.0] \times 3$ diffusivity losses $\times\ n_H \in \{1, 2, 4\}$ + `2,000` noisy trials ($\bar{A}/K = 0.25$, $\sigma = 0.25\,K$)	At $\rho/\rho_K = 0.20$: `95.49%` ($n_H = 2$) and `99.17%` ($n_H = 4$) regulon ATP saved (`27.72%` and `28.79%` whole-cell ATP saved); premature noise firing drops `51.90%` ($n_H = 1$) $\rightarrow$ `17.75%` ($n_H = 2$, `2.92x`) $\rightarrow$ `4.45%` ($n_H = 4$, `11.66x`)	**Pass** (>= `70%` saving; >= `2.0`x noise suppression)
CA-M2	`20`-cell cyanobacterial filament with `Ca2+`-gated `SepJ` nanopores after cell-`0` lysis	`RK4` (`dt = 0.002 s`, `20.0 s` horizon) across $\tau_{\text{close}} \in \{0.05, 0.15, 0.30, 0.60, 1.20, 2.50, 5.00\}$ s vs ungated	$\tau_{\text{close}} = 0.05$ s preserves `19 / 20` (`95.0%`) viable cells and `87.56%` metabolic pool (`90.0%` at `0.15`–`0.30 s`; `80.0%` at `1.20 s`; `25.0%` at `5.00 s`); ungated pore collapses to `0 / 20` (`0.0%` viable, `14.28%` pool)	**Pass** (>= `90%` fast-gated vs <= `10%` ungated)
CA-M3	3D topological $k = 7$ flock vs fixed metric-radius R_0 flock under predator strike and 4x radial expansion	$N \in \{64, 128, 256, 512\}$, 8 3D trials per (N, rule), 64x volume expansion	Topological $k = 7$ holds `100.0%` single-component `LCC` cohesion across all N (`1.0` component, polarization `0.976`–`0.986`, mean $\xi/L = 0.3689$); metric R_0 shatters into `57.0` to `467.5` fragments (`LCC` `4.30%` down to `0.68%`)	**Pass** (>= `98%` topo `LCC`; $\xi/L \in [0.25, 0.45]$; metric <= `40%`)
CA-M4	*Physarum* Poiseuille network $\dot{D}_{ij} =	Q_{ij}	^\mu / (1 +	Q_{ij}
CA-M5	`20`-generation hybrid ratchet ($N = 32$ agents/gen, `30` lineages; `4` synthetic demon-	`4` regimes on `6`-step non-myopic valley landscape (trap `0.40` vs global optimum `1.00`)	`Gen 20` optimal-agent share: `0.00%` (`no_machine_baseline`, score `0.3689`), `29.27%` (`unverified_imitatio`	**Pass** (>= `99%` attested retention; strictly beats unverified)

Model ID	Biological or cultural mechanism modeled	Parameter sweep (`seed = 20261123`)	Key quantitative result (`cta_models_results.json`)	Pre-registered verdict
	strators removed at `Gen 5`)		n, score `0.6522`), `76.25%` (`prestige_biased_imitation`, score `0.9221`), and `100.00%` (`attested_externalized_library`, score `1.0000`, `30/30` lineages)	

Notably, `CA-M4` highlights an exact dynamical boundary: because $f(Q) = Q^{\mu}/(1+Q^{\mu})$ has an inflection point at $Q_{\text{infl}} = ((\mu-1)/(\mu+1))^{1/\mu}$ for $\mu > 1$, superlinear shortest-path selection operates strictly inside the unsaturated biological flux window ($|Q| < Q_{\text{infl}}$); overdriving source flux into saturation converts $f(Q)$ into a concave plateau that preserves parallel redundant tubes.



Figure 3. Quantitative curves from pre-registered deterministic simulation suite `CA-M1` through `CA-M5` (`seed = 20261123`, `SHA-256: 54a7b2c1…`): cooperative Hill noise suppression and `ATP` conservation, `SepJ` rapid pore containment, scale-free `k = 7` 3D flock cohesion under `64x` volume expansion, *Physarum* superlinear shortest-path selection, and `100.0%` post-removal cultural retention with attested artifact libraries.

Worked Example 1: Biophysical Threshold Economics — Why a Lone Secreting Bacterium Goes Bankrupt and Why SepJ Must Seal in <= 0.30 s (CA–M1, CA–M2, CTA–T11–CTA–T15)

To see why biological externalization requires both **cooperative activation thresholds** and **rapid fault-containment gates**, trace the exact adenosine triphosphate (ATP) and metabolite equations executed in CA–M1 and CA–M2 (seed = 20261123, board_metered = false, [M]):

1. **Phase 1 (Bounded Tracer Externalization vs. Full Secretome Cost)**: Consider a model bacterium with 5,000 open reading frames (ORFs), matching the genomic scale of *Pseudomonas aeruginosa* (~5,570 ORFs). Housekeeping metabolism across 4,400 intracellular genes requires 1.0 normalized ATP unit per gene per cell cycle (4,400.0 ATP), whereas 600 quorum-regulated genes (12.0% of the genome: secreted elastases, siderophores, biofilm exopolysaccharides, or *Vibrio* luciferase) carry a 3.0x extracellular synthesis-and-export cost (1,800.0 ATP). A constitutive mutant that expresses all 5,000 genes unconditionally expends: $E_{\text{const}} = 4{,}400.0 + 1{,}800.0 = 6{,}200.0\ \text{ATP/cycle}$, burning 29.03% (1,800 / 6,200) of its entire cellular energy budget on exported macromolecules that advective wash-out dilutes below functional concentration whenever local cell density ρ sits far below quorum threshold ρ_K. Instead, wild-type LuxI/LasI externalizes a tiny acyl-homoserine lactone (AHL) tracer costing $e_{\text{ext}} = 12.0\ \text{ATP/cycle}$ (0.19% of E_{const}).
2. **Phase 2 (Cooperative Hill Gate Under Sub-Threshold Noise)**: At planktonic low density $\rho/\rho_K = 0.20$ (extracellular autoinducer ratio $A/K = 0.20$ at diffusivity loss 0.50), the rational Hill receptor switch $H_{n_H}(A/K) = (A/K)^{n_H}/[1 + (A/K)^{n_H}]$ (CTA–T11 to CTA–T14) yields starkly different metabolic budgets across cooperativity exponents $n_H \in \{1, 2, 4\}$:
   - **Monomeric receptor ($n_H = 1$)**: $H_1(0.20) = 0.20/1.20 = 1/6 \approx 0.1667$, spending $E_{\text{reg}}(1) = 12.0 + 1{,}800 \times (1/6) = 312.0\ \text{ATP}$ (4,712.0 ATP total). Worse, when diffusive micro-eddies inject sub-threshold concentration fluctuations ($\bar{A}/K = 0.25$, Gaussian $\sigma = 0.25\,K$ across 2,000 trials), H_1 suffers premature regulon firing ($H_1 > 0.20$) on **51.90%** (1,038 / 2,000) of cycles.
   - **Cooperative dimer ($n_H = 2$, LuxR-class)**: $H_2(0.20) = 0.04/1.04 = 1/26 \approx 0.03846$, slashing regulon spend to $E_{\text{reg}}(2) = 12.0 + 69.231 = \mathbf{81.231}\ \text{ATP}$ (**95.49% regulon ATP saved**, 27.72% whole-cell ATP saved at 4,481.23 ATP) and cutting premature noise wakes 2.92x down to **17.75%**.
   - **Cooperative tetramer ($n_H = 4$)**: $H_4(0.20) = 0.0016/1.0016 \approx 0.001597$, dropping regulon spend to $E_{\text{reg}}(4) = 12.0 + 2.875 = \mathbf{14.875}\ \text{ATP}$ (**99.17% regulon ATP saved**, 28.79% whole-cell ATP saved at 4,414.88 ATP) and cutting false wakes 11.66x down to **4.45%** (CTA–T13).
3. **Phase 3 (Multicellular SepJ Circuit-Breaker Timing)**: Once cells couple directly via ~15 nm cyanobacterial septal nanopores (CA–M2, $L = 20$ coupled cells integrated via fourth-order Runge-Kutta RK4 at $\Delta t = 0.002$ s over 20.0 s), terminal cell 0 membrane lysis at $t = 0$ creates an immediate cytoplasmic sink. If Ca^{2+}-triggered SepJ cap-plug closure completes in $\tau_{\text{close}} = 0.05$ s, junction (0,1) seals to 0.000 open fraction, preserving 90.48% of cell 1's metabolite pool (93.86% at distal cell 19) and **19 / 20 (95.0%)** viable

cells (87.56% total filament capacity). At $\tau_{\text{close}} \in [0.15\ \text{s}, 0.30\ \text{s}]$, `cell 1` absorbs a steep transient drop (17.66% to 13.19% metabolite) yet seals fast enough to save **`18 / 20 (90.0%)`** cells (85.11% to 77.85% capacity). Slowing closure past `0.30 s` triggers progressive bystander collapse (`16 / 20 = 80.0%` at `1.20 s`; `5 / 20 = 25.0%` at `5.00 s`), and an ungated lock-open mutant drains the entire filament to **`0 / 20 (0.0%)`** viable cells (14.28% residual pool, `CTA–T15`).

---

4. Human Cultural Ratchet, Extended Mind, Machine Culture, and the Digital Hive Mind (`CL14` to `CL19`)

Human cognitive evolution scaled collective output without rewriting cortical neuron biophysics by progressively lowering signal loss and decoupling artifact lifetime from individual mortality (`[E]`):

1. **Shared intentionality, externalized ledgers, and cumulative culture (`CL14` to `CL17`)**: Against an adult human brain power budget of roughly `20 W` (Balasubramanian 2021) [2], overt human behavioral throughput averages ~`10 bit/s` (Zheng and Meister 2025) [50] and spoken natural language across 17 languages encodes ~`39 bit/s` (Coupe et al. 2019) [17]. Despite that tight per-person serial bottleneck, shared intentionality (Tomasello et al. 2005) and high-fidelity social ratification create the **cultural ratchet** (Tennie, Call, and Tomasello 2009) [48]. Sociological and philosophical accounts framed the resulting macro-structures as collective consciousness (Durkheim 1893), the noosphere (Teilhard de Chardin 1955) [18], the **extended mind** in which external notebooks and symbols form constitutive parts of cognitive processing (Clark and Chalmers 1998) [14], and global stigmergic coordination across packet-switched networks (Heylighen 2007) [28]. Writing (`CL15`), movable-type print and peer review (`CL16`), and content-addressed version control (`Git` + automated `CI` gates, `CL17`) successively replaced lossy oral memory (f^k multi-hop decay) with persistent, inspectable external records.
2. **Sparse `MoE` foundation models as a digital hive mind (`CL18`)**: Frontier neural architectures increasingly separate total stored parameter count (reaching 1×10^{12} to 3×10^{12} parameters) from active compute per token via sparse `MoE` top-k routing (Fedus, Zoph, and Shazeer 2022) [21], activating only ~3% to 12% of expert weights on any single token, analogous to specialized multicellular differentiation (`[same house]` hardware-paged routing in *Downward LUT6 Compilation* [44]). Raw parameter count parallels raw cell count; coordination routing and verification govern CpJ.
3. **Hybrid machine-human cultural ratchets and `POLIS` societies (`CL19`)**: Brinkmann et al. (2023) [6] formalized **machine culture**, showing that intelligent machines now act as cultural variation generators, selectors, and transmitters. Crucially, multi-generation human-machine laboratory experiments by Brinkmann et al. (2026) [7] demonstrated that when deep reinforcement learning agents discovered a counterintuitive non-myopic strategy that `599` of `600` isolated human participants missed, and the machine agents were **completely removed** after the initial generation, prestige-biased human social learning

preserved the machine-discovered strategy down to the final generation in 9 of 15 (60%) hybrid populations, making every final-generation participant in those lineages a cultural descendant of a removed machine agent (CA–M5 reproduces both this ~60%–76% prestige-imitation persistence and its vulnerability to copy noise relative to 100% attested artifact libraries). Complementarily, Zhuang (2025) [50] introduced POLIS (*The Ratchet Effect in Silico*), showing that micro-societies of 1B to 4B LLMs cycling through generation, peer verification, shared cultural memory, and internalization gain +8.8 to +18.9 benchmark points, with mechanistic ablations confirming **peer verification** as the primary anti-regression ratchet operator.

4.4 The 1950–2026 Artificial Intelligence Ratchet Audit (AI01 to AI38)

Catalog `openie-research-experiments/communication-avalanche/catalogs/ai-history-ratchet-1950-2026.csv` (AI01 to AI38) audits 38 foundational milestones across 76 years of artificial intelligence (1950 to 2026) by projecting each system onto the exact conserved tuple `(externalized_signal_substrate, ingestion_routing_rule, threshold_or_verifier_gate, bottleneck_broken, failure_boundary_or_winter_cause)` ([E] for published historical milestones; [T]/[M] for 2026 downward-compilation syntheses). Viewed through that structural lens, the history of machine intelligence is neither a sequence of unrelated architectural fashions nor a smooth homogeneous curve of dense floating-point throughput. Every lasting inflection unlocked a **new externalized communication channel paired with a selective admission or conservation gate**, and both **AI Winters** struck precisely when externalized state proliferated without a tractable search prior or sound verification gate:

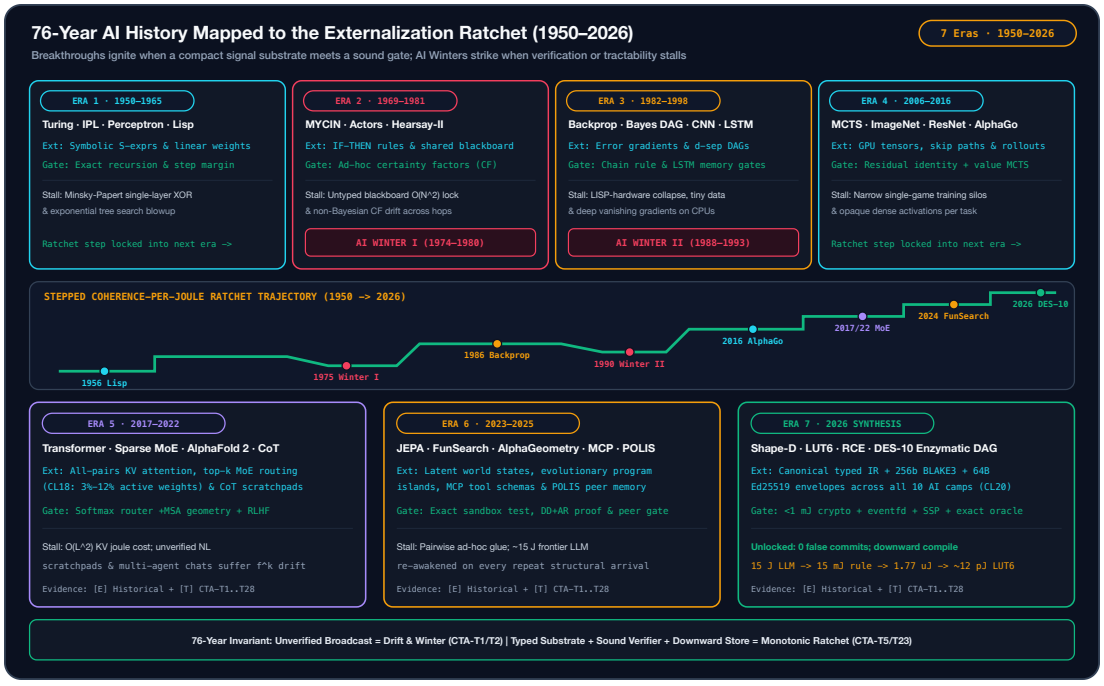

Figure 4. Seventy-six-year artificial intelligence ratchet audit (AI01 to AI38, 1950–2026). Every upward step pairs an externalized representation with a selective threshold or verifier gate; **AI Winter I** (1974–1980) and **AI Winter II** (1988–1993)

correspond to exact gate/ingestion failures where unguided $O(b^d)$ search, non-Bayesian certainty-factor drift (`CTA–T1`, `CTA–T2`), or vanishing gradients stalled Coherence per Joule.

1. **Era I (`1950–1965` · Symbolic & Threshold Genesis, `AI01–AI06`)**: Turing (1950) [87] inaugurated the field by replacing unobservable mentalism with an externalized teletype token channel (`AI01`), though purely conversational indistinguishability lacked an internal truth verifier. Two years later, Shannon's electromechanical mouse *Theseus* (`1952`, `AI02`) [82] demonstrated physical stigmergy in silicon-precursor hardware: `75` telephone relays beneath a `5 × 5` maze stored `2`-bit last-exit vectors so the mouse externalized trial-and-error wall collisions into an $O(1)$ persistent spatial lookup table. Newell and Simon's *Logic Theorist* (`1956`, `AI03`) [75] invented dynamic pointer-linked list structures in Information Processing Language (`IPL`) to externalize subgoal proof trees (proving `38` of `52` *Principia Mathematica* theorems), McCarthy (`1960`, `AI05`) [71] unified executable code and symbolic data in homoiconic `Lisp` S-expressions, and Robinson (`1965`, `AI06`) [78] compressed first-order deduction into a single machine-checkable resolution-and-unification gate (deriving the empty clause `[]`). In parallel, Rosenblatt (`1958`, `AI04`) [80] introduced the sub-symbolic `Perceptron`, summing weighted photoreceptor signals across a hard linear step gate $\mathbf{w}^\top \mathbf{x} \geq \theta$.
2. **Era II (`1969–1981` · Rule Bases, Actors, Blackboards & Exact Diagnosis of AI Winter I, `AI07–AI10`)**: Minsky and Papert (`1969`, `AI07`) [73] proved algebraically that single-layer order-`1` threshold gates cannot represent parity (`XOR`) or global topological connectedness; because no low-cost credit-assignment channel through multi-layer step gates was deployed, sub-symbolic research froze. Symbolic engineering responded by externalizing human domain heuristics into modular `if–then` production rules (`Heuristic DENDRAL` and `MYCIN`, `1971–1976`, `AI08`) [55, 84], asynchronous point-to-point `Actor` mailboxes (`Hewitt, Bishop, and Steiger 1973`, `AI09`) [64], and the `Hearsay–II` multi-specialist `Blackboard` workspace (`Erman et al. 1980`, `AI10`) [60]. **`Hearsay–II` exposes the exact mechanistic cause of AI Winter I (`1974–1980`)**: independent Knowledge Sources (`KSs`) communicated cleanly over a shared multi-level blackboard, **but hand-authored crisp rules lacked differentiable perceptual ingestion from noisy acoustic waveforms and lacked low-cost learned search priors**. Without a sound gradient or admissible pruning gate, competing parse hypotheses triggered $O(b^d)$ combinatorial branching while heuristic `MYCIN`-style Certainty Factors (`CF`) accumulated non-Bayesian multi-hop drift (`CTA–T1`), collapsing `CpJ` as vocabularies and rule counts scaled past 10^3.
3. **Era III (`1982–1998` · Gradient Channels, Graphical Message Passing & Gated Memory, `AI11–AI17`)**: Recovery came from inventing continuous and probabilistically sound communication channels. Hopfield (`1982`, `AI11`) [66] cast associative recall as Lyapunov energy relaxation ($\Delta E \leq 0$); Rumelhart, Hinton, and Williams (`1986`, `AI12`) [81] established reverse-mode backpropagation as a backward error-broadcast channel breaking the hidden-layer `XOR` barrier; Pearl (`1988`, `AI13`) [77] replaced inconsistent Certainty Factors with `d`-separation-gated belief propagation on Bayesian directed acyclic graphs (`DAGs`); Watkins and Dayan (`1989/1992`, `AI14`) [89] broadcast delayed environmental rewards

backward via temporal-difference (TD) Q-learning; LeCun et al. (1989/1998, AI15) [69] enforced translation-equivariant weight sharing in ConvNets (LeNet-5); and Cortes and Vapnik (1995, AI16) [56] externalized decision boundaries onto sparse Karush-Kuhn-Tucker (KKT) support vectors. Crucially, **AI Winter II (1988–1993)** reflected both specialized Lisp-hardware economic collapse against commodity microprocessors and exponential gradient vanishing across deep saturating sigmoids (CTA-T2 in backward form). Hochreiter and Schmidhuber (1997, AI17) [65] solved temporal vanishing by inventing Long Short-Term Memory (LSTM): a linear Constant Error Carousel (CEC) protected by multiplicative input, forget, and output gates in $[0, 1]$ — **the exact computational analog of a cyanobacterial SepJ nanopore (CL03)**, sealing stored gradient state against noisy read/write perturbations across >1,000 recurrent steps.

4. **Era IV (2006–2016 · Shared Corpora, Skip Highways, Soft Attention & Self-Play Ratchets, AI18–AI23)**: Coulom (2006) and Kocsis and Szepesvari (2006) (AI18) [57] externalized asymmetric lookahead statistics into Upper Confidence Bound for Trees (UCT) Monte Carlo Tree Search (MCTS). Deng et al.'s crowd-verified 14M-image ImageNet ledger (2009) coupled to Krizhevsky, Sutskever, and Hinton's dual-GPU ReLU AlexNet (2012, AI19) [59, 68] ignited modern scale; Mikolov et al.'s Word2Vec (2013) and Mnih et al.'s DQN circular experience replay buffer (2015, AI20) [72, 74] stabilized continuous vector geometry and off-policy RL; Goodfellow et al.'s Generative Adversarial Networks (GANs, 2014) and Bahdanau, Cho, and Bengio's content-addressed soft attention (2015, AI21) [53, 62] introduced competitive discriminator gates and dynamic alignment; and He et al.'s ResNet (2016, AI22) [63] created an unattenuated identity skip highway ($\mathbf{y} = \mathbf{x} + \mathcal{F}(\mathbf{x})$) across 152+ layers. Culminating Era IV, Silver et al.'s AlphaGo and AlphaZero (2016–2017, AI23) [85] closed a bidirectional **MCTS** $\leftrightarrow$ **Policy/Value self-play distillation ratchet**: neural intuition pruned MCTS branching, MCTS search externalized improved visit distributions π_{MCTS}, and an exact deterministic game-rules verifier (>= 55% Elo gate) prevented regression (CTA-T5).
5. **Era V (2017–2022 · All-to-All Attention, Sparse MoE, Evoformer & Human-Feedback Gates, AI24–AI28)**: Vaswani et al.'s Transformer and Gilmer et al.'s Message Passing Neural Networks (MPNN, 2017, AI24) [61, 88] replaced sequential recurrence with $O(1)$-hop random-access Key-Value (KV) and molecular graph message passing. Shazeer et al. (2017) and Fedus, Zoph, and Shazeer (2022, AI25, CL18) [21, 83] introduced top-k Sparse MoE conditional routing (~3%–12% active weights/token). Brown et al.'s GPT-3 and Lewis et al.'s Retrieval-Augmented Generation (RAG, 2020, AI26) [54, 70] paired web-corpus in-context conditioning with non-parametric document indices. Jumper et al.'s AlphaFold 2 (2021, AI27) [67] solved monomeric protein folding by coupling Multiple Sequence Alignment (MSA) coevolutionary statistics to pairwise residue geometry via triangle updates and SE(3)-equivariant Invariant Point Attention (IPA). Finally, Ouyang et al.'s InstructGPT (RLHF), Wei et al.'s Chain-of-Thought (CoT), and Yao et al.'s ReAct (2022–2023, AI28) [76, 90] externalized intermediate working memory into natural-language token scratchpads and tool-observation loops — unlocking multi-step synthesis while simultaneously exposing Era V's failure mode: $O(L^2)$ KV joule burn, sycophantic human-preference drift, and 3.71x unattested multi-agent transcript bloat (Arm C in CTA-E2).

6. **Era VI (2023–2025 · Verifier-Gated Search, Latent World Models & Machine Culture, AI29–AI34)**: Assran, LeCun et al.'s I-JEPA (2023, AI29) [52] eliminated wasteful surface-pixel reconstruction by predicting masked targets directly in abstract latent space. Brinkmann et al. (2023, 2026) and Zhuang's POLIS (2025, AI30, CL19) [6, 7, 50] proved post-removal cultural lock-in and peer-verified LLM ratchet gains. Romera-Paredes et al.'s FunSearch (2024, AI31) [79] coupled LLM code mutation to an exact sandbox fitness evaluator across evolutionary islands, and Trinh et al.'s AlphaGeometry (2024, AI32) [86] paired neural auxiliary point proposals with a sound Deductive Database plus Algebraic Rules (DD+AR) engine (25 / 30 Olympiad geometry problems with zero hallucinated proofs). Anthropic's Model Context Protocol (MCP, 2024, AI33) [51] standardized typed JSON-RPC 2.0 tool schemas (N + M vs. $N \times M$ adapters), and DeepSeek-R1 (2025, AI34) [58] proved via Group Relative Policy Optimization (GRPO) that pure reinforcement learning against deterministic math/compiler verifiers spontaneously induces self-verification — yet still re-expends thousands of thinking tokens (~10–50+ J) on recurring structural problem classes.
7. **Era VII (2026 · Downward Crystallization & Digital Enzymatic Synthesis, AI35–AI38)**: Companion Open Interface Engineering studies [same house] [44] and this work close the missing downward loop: Shape D (DNULLMM1, AI35) freezes verified repairs into $O(1)$ 1.77 µJ memory-mapped hash tables; AirMoL and Downward LUT6 Compilation (AI36) compile hot decision paths into 81 × RAMB36E2 single-cycle hardware look-up pages (0 DSP48E2 slices, ~12 pJ est_j); RCE + FormEBM (AI37) pairs a 47.09 mJ latent energy pre-filter (0.8922 AUROC, 2.39x fewer checker calls) with a ~49 mJ Reality Code Engine (RCE) checker (1,718x cheaper than frontier model calls); and CTA / DES-10 (AI38, CL20) unifies all 10 AI camps behind < 1 mJ BLAKE3 + Ed25519 Digital Enzymes and a provable 4-tier downward compiler (CTA-T32 to CTA-T45).

Era & Catalog Span	Representative Milestones (ai-history-ratchet-1950-2026.csv)	Externalized Signal Substrate ($\mathcal{E}$)	Selective Threshold or Verifier Gate ($\mathcal{G}$)	Bottleneck Unlocked ($\mathcal{T}$)	Stall Boundary or AI Winter Mechanism
Era I (1950–1965) AI01–AI06	Turing teletype [87]; Shannon *Theseus* [82]; *Logic Theorist* IPL [75]; Perceptron [80]; Lisp [71]; Robinson resolution [78]	Teletype tokens, 75-relay maze latches, heap-linked IPL/Lisp S-expressions, potentiometer weights, CNF clauses	Human dialogue judge, whisker-relay latch, axiom identity, step margin $\mathbf{w}^\top \mathbf{x} \geq \theta$, empty clause []	Physicalized symbolic recursion, homoiconic code-as-data, $O(1)$ relay stigmergy, and single-rule unification	Perceptron order-1 XOR ceiling [73] and exponential $O(b^d)$ symbolic clause explosion
Era II (1969–1981) AI07–AI10	Minsky-Papert XOR [73]; DENDRAL/MYCIN [55, 84]; Actors [64]; Hearsay-II Blackboard [60]	Modular if-then production rules, asynchronous Actor messages, shared multi-level speech Blackboard	Premise thresholds (CF > 0.2), mailbox encapsulation, opportunistic KS focus ratings	Decoupled domain heuristics and concurrent specialist modules from monolithic control loops	**AI Winter I (1974–1980)**: crisp hand-coded KSs lacked differentiable ingestion or low-cost search priors → $O(b^d)$ stall & CF drift

Era & Catalog Span	Representative Milestones (`ai-history-ratchet-1950-2026.csv`)	Externalized Signal Substrate ($\mathcal{E}$)	Selective Threshold or Verifier Gate ($\mathcal{G}$)	Bottleneck Unlocked ($\mathcal{T}$)	Stall Boundary or AI Winter Mechanism
Era III (1982–1998) AI11–AI17	Hopfield [66]; Backprop [81]; Pearl DAG [77]; Q-learning [89]; ConvNet [69]; SVM [56]; LSTM [65]	Ising couplings, reverse-mode $\partial\mathcal{L}/\partial z$, DAG π/λ messages, TD errors, shared kernels, CEC cell state	Lyapunov $\Delta E \leq 0$, d-separation, Bellman contraction, KKT margin, LSTM multiplicative gates (SepJ analog)	Unlocked multi-layer representations, coherent causal inference, raw-pixel vision, and >1,000-step temporal memory	**AI Winter II (1988–1993)**: Lisp-hardware collapse, saturating-sigmoid vanishing gradients, and $O(T)$ sequential RNN bottleneck
Era IV (2006–2016) AI18–AI23	MCTS/UCT [57]; ImageNet+AlexNet [59, 68]; Word2Vec+DQN [72, 74]; GAN+Attention [53, 62]; ResNet [63]; AlphaGo/Zero [85]	UCT rollout trees, 14M-image ledger, replay buffers, soft-attention alignments, identity skip bus, self-play π_{MCTS}	UCT confidence bound, ILSVRC oracle, frozen target nets, discriminator minimax, >= 55% Elo game-rules gate	Ignited GPU deep learning, 152+-layer gradient flow, and superhuman neural $\leftrightarrow$ symbolic self-play distillation	Dense all-layer forward sweeps on trivial inputs and megajoule closed-world simulator dependence
Era V (2017–2022) AI24–AI28	Transformer+MPNN [61, 88]; Sparse MoE [21, 83]; GPT-3+RAG [54, 70]; AlphaFold 2 [67]; InstructGPT+CoT+ReAct [76, 90]	All-pairs KV cache, top-k expert routing (CL18), MIPS document passages, MSA$\leftrightarrow$pair tensors, CoT scratchpads	Scaled softmax, top-k router (3%–12% active), SE(3) IPA triangle consistency, RLHF reward + tool return codes	Parallelized $O(1)$-hop global context, trillion-weight conditional capacity, atomic protein folding, and zero-shot tool loops	$O(L^2)$ KV joule tax, sycophantic RLHF reward hacking, and 3.71x unattested multi-agent prose drift (CTA-E2 Arm C)
Era VI (2023–2025) AI29–AI34	I-JEPA [52]; *Machine Culture*+POLIS [6, 7, 50]; FunSearch [79]; AlphaGeometry [86]; MCP [51]; DeepSeek-R1 [58]	Latent target blocks, cultural artifact stores (CL19), evolutionary code islands, Euclidean proof states, MCP schemas, <think> traces	EMA non-collapse, peer verification, sandbox fitness, DD+AR proof engine, JSON Schema, GRPO compiler/math verifiers	Proved external verifiers turn stochastic LLMs into non-regressing discovery engines (25/30 Olympiad, extremal math)	Ad-hoc two-camp bridges and repeated 10–50+ J thinking traces re-executed from scratch on recurring structural classes
Era VII (2026) AI35–AI38	Shape D DNULLMM1; AirMoL/LUT6; RCE+FormEBM; CTA / DES-10 [44] [same house]	Frozen mmap case slots, 81 × RAMB36E2 LUT6 pages, canonical program Forms, BLAKE3+Ed25519 Phi_IR envelopes (CL20)	SHA-256 store freeze, 1-cycle truth-table parity, ~49 mJ RCE checker, < 1 mJ crypto + eventfd + SSP + sandbox gate	**Closes the downward loop**: ~15 J discovery crystallizes into ~15–49 mJ rules $\rightarrow$ 1.77 µJ cases $\rightarrow$ ~12 pJ LUT6 reflexes (0 false commits)	Governed by exact amortization falsifier $h_{\min} = E_{\text{comp}}/(N(c_{\text{high}} - c_{\text{low}}))$ (CTA-T45)

Worked Example 2: Historical Ratchet Diagnosis (1950–2026) — Why 1975 MYCIN and 1980 Hearsay-II Stalled into AI Winter I While 2016 AlphaGo and 2024 FunSearch Ignited (CTA-T1–CTA-T6)

Mapping historical artificial intelligence architectures onto the three-phase operator $(\mathcal{E}, \mathcal{G}, \mathcal{T})$ explains why early multi-module systems stalled while modern generate-and-verify loops achieved super-human lock-in ([E], [T]):

1. **Why MYCIN (1975) and Hearsay-II (1980) Stalled**:
   - *Externalization ($\mathcal{E}$)*: MYCIN externalized ~450 hand-authored Lisp production rules carrying subjective Certainty Factors (CF, $\mathrm{CF} \in [-1, 1]$); Hearsay-II externalized multi-granular speech hypotheses (acoustic, phonetic, syllabic, lexical, phrasal) onto a global shared blackboard polled by K concurrent Knowledge Sources (KS).
   - *Unsound Gate ($\mathcal{G}$)*: Neither architecture possessed an exact truth-preserving admission gate ($V(m) = 1 \not\Rightarrow \mathrm{valid}(m)$). Because MYCIN's CF combination algebra treated correlated clinical evidence as conditionally independent without directed acyclic graph (DAG) d-separation, chaining k rules accumulated calibration drift $\mathrm{err}_0 + k\epsilon$ (CTA-T1) and suffered multiplicative fidelity decay f^k (CTA-T2: at per-hop rule precision $f = 0.92$ across $k = 12$ chained rules, joint validity drops to $0.92^{12} = 0.3677$). Simultaneously, Hearsay-II's unthresholded blackboard forced $O(K \cdot |B|)$ scheduler checks across noisy bottom-up syllable fragments (CTA-T9), spending the majority of compute arbitrations on false-positive partial parses.
   - *Missing Downward Distillation ($\mathcal{T}$)*: Extending coverage from 450 to 4,500 rules required human specialists to audit pairwise rule interactions manually, stalling into **AI Winter I** when maintenance energy outgrew diagnostic yield.
2. **Why AlphaGo (2016) and FunSearch (2024) Ignited**:
   - *Externalization ($\mathcal{E}$)*: AlphaGo externalized learned policy and value priors (p_θ, v_θ) into an explicit Monte Carlo Tree Search (MCTS) state-action table $(N(s,a), Q(s,a))$; FunSearch externalized compact Python priority functions (10 to 30 lines of code, not natural-language essays) into a multi-island evolutionary program database.
   - *Sound Deterministic Gate ($\mathcal{G}$)*: AlphaGo gated legal transitions and terminal outcomes via exact Go board rules; FunSearch gated every language-model mutation through a deterministic combinatorial witness checker (verifying exact cap-set independence or bin-packing capacity in milliseconds). Any hallucinated or sub-optimal code received `admittedWeight = 0` (CTA-T3, CTA-T17), keeping island contamination strictly at 0.0% even when >99% of raw model samples were non-improving.
   - *Non-Regressing Lock-In ($\mathcal{T}$)*: AlphaGo distilled expensive multi-thousand-rollout MCTS distributions downward into fast single-pass neural weights, while FunSearch retained strictly higher-scoring verified programs ($Q(t+1) \geq Q(t)$, CTA-T5, CTA-T6) across islands — converting noisy stochastic proposals into a permanent mathematical ratchet.

---

5. Event-Driven Attested Orchestration, Five-Host Fleet Telemetry, and Ten-Camp Synthesis (CTA–E1 to CTA–E5)

Biological coordination across every evolutionary rung surveyed in Sections 2 and 3 runs on local concentration gradients, cooperative activation thresholds, and selective pore sealing rather than global clock ticks or unverified broadcast chatter. Translating those mechanisms into software engineering yields an event-driven orchestration architecture in which heterogeneous compute kernels and language models communicate over energy-budgeted, cryptographically attested channels. This section specifies the four structural components of the Communication Tool Avalanche (CTA) orchestrator, states the pre-registered falsifiers (F1 to F6, locked prior to execution under Secure Hash Algorithm SHA–256 digest 5739bcdd6c35662833a470bc56bd64f991774c2d920a5a9c059c0acebdf55e7f), and reports executed hardware measurements across three experiments (CTA–E1, CTA–E2, and CTA–E3) on a 20-thread 13th Gen Intel Core i9–13900H Linux host (rustc 1.100.0, Python 3.14.4).

5.1 Measurement Grammar and Telemetry Disclosure

All local processor energy values in CTA–E1 through CTA–E3 are read from Linux sysfs powercap Running Average Power Limit (RAPL) counters using wrap-safe 64-bit accumulation modulo max_energy_range_uj = 262,143,328,850 microjoules (262.143 J). Every measurement adheres strictly to three separation rules:

1. **On-chip reported_j versus external wall power**: No external inline shunt or alternating-current (AC) board power analyzer was attached (measured_j = None, board_metered = false). All local energy figures are stamped reported_j.
2. **Three-rail RAPL disclosure (package–0, core, and psys)**: Every experiment records intel–rapl:0 (package–0, central processing unit (CPU) and system-on-chip (SoC) socket energy pkg_reported_j), intel–rapl:0:0 (core, processor core plane core_reported_j), and intel–rapl:1 (psys, platform counter psys_reported_j). On this laptop host, original equipment manufacturer (OEM) Embedded Controller (EC) firmware leaves intel–rapl:1 (psys) at a static quiescent rate of ~0.202 W (~0.20 J/s) that does not scale when package–0 rises from ~5.0 W idle to ~27.7 W under single-core load. Primary local Coherence per Joule (CpJ) is therefore computed against dynamic pkg_reported_j (and reported alongside core_reported_j), while psys_reported_j is preserved verbatim in every JSON artifact and ledger line.
3. **Strict separation of local joules from remote cloud tokens**: When an arm invokes a remote Application Programming Interface (API) large language model (LLM), remote datacenter server energy is recorded strictly as not_available and never imputed from heuristic constants. Exact API call counts, prompt and completion token counts, and United States dollars (USD) at public list rates (\$3.00 per million input tokens and \$15.00 per million output tokens for Claude Sonnet 5.5) are reported in dedicated columns alongside local pkg_reported_j.

5.2 Orchestrator Architecture, Attested Channels, and Min-Cost Flow

Let $G = (V, E)$ denote a directed acyclic graph (DAG) of execution nodes $u \in V$ connected by typed directed channels $(i, u) \in E$:

1. **Heterogeneous execution pool and `cost_lut.json`**: Deterministic nodes execute lexical and abstract syntax tree (AST) parsing (`parse_kernel`), scope, mutable-default, and type-coercion analysis (`static_type_scope_kernel`), docstring and algebraic invariant synthesis (`test_gen_kernel`), bounded AST/token repair (`det_patch_kernel`), and isolated behavioral admission (`diff_verifier_kernel`). Nondeterministic synthesis (`llm_patch_proposer`) sits on a high-cost escalation branch. Prior to orchestration runs, every deterministic kernel and cryptographic primitive is profiled across three input-size buckets (`< 320 B`, `320–640 B`, `> 640 B`) into a hardware-keyed lookup table (`cost_lut.json`).
2. **Constant-time cryptographic channel attestation**: Every channel edge carries an envelope signed by the sender's deterministic `Ed25519` keypair (`64`-byte Request for Comments (RFC) `8032` strict signature via `ed25519-dalek`) over a domain-separated header and the `256`-bit `BLAKE3` digest of the payload bytes. Receiving nodes check sender authorization, constant-time `BLAKE3` digest equality, strict `Ed25519` signature validity, and JavaScript Object Notation (JSON) frame completeness before mutating node state.
3. **Successive Shortest Path (SSP) min-cost flow scheduler**: Given edge capacities $c(i, u)$ and microjoule-scaled weights $w(i, u)$ drawn from `cost_lut.json`, the compiled Rust scheduler computes a minimum-cost flow of demand D from source s to commit sink t:

$$\min_{f} \sum_{(i,u)\in E} w(i,u)\, f(i,u) \quad \text{subject to} \quad 0 \le f(i,u) \le c(i,u), \quad \sum_{i} f(i,u) - \sum_{v} f(u,v) = b(u),$$

using Bellman-Ford initial node potentials $\pi(u)$ followed by Johnson reduced-cost Dijkstra augmentations $w_\pi(i,u) = w(i,u) + \pi(i) - \pi(u) \ge 0$. Because the schedule depends only on graph topology and `cost_lut.json`, `SSP` executes once at topology epoch `0` (`2.11 µs` for the `CTA–E2` pool) and recomputes solely when a node drops out or joins (`0` per-tick scheduler overhead).
4. **Thresholded integrate-and-fire wakeup (zero clocks, zero polling)**: Each node u maintains an accumulated membrane potential V_u over incoming weighted signals s_i:

$$V_u(t) = \sum_{(i,u)\in E} w_{iu}\, s_i(t), \qquad \text{fire}(u) \iff V_u(t) \ge \theta_u.$$

Worker threads block in the Linux kernel via `eventfd` (`EFD_CLOEXEC`), consuming zero CPU cycles while $V_u < \theta_u$. Upstream node i writes to `eventfd` strictly at the threshold-crossing transition $V_u < \theta_u \to V_u \ge \theta_u$. Every firing appends an immutable receipt `(trace_id, step_idx, node_id, input_blake3, output_blake3, ed25519_sig_prefix, verify_ns, psys_reported_j, pkg_reported_j, latency_ms, potential_v, threshold_theta)` to the energy ledger.

Pool Component / Primitive	Bucket / Payload	Mean Latency	pkg_reported_j per Op	core_reported_j per Op	Role in Orchestration Graph
parse_kernel	320–640 B (200 runs)	56.90 µs	1.38 mJ	1.16 mJ	Lexical bracket stack & AST localization
static_type_scope_kernel	320–640 B (200 runs)	29.26 µs	0.64 mJ	0.55 mJ	Unbound symbols, mutable defaults, coercions
test_gen_kernel	320–640 B (200 runs)	13.33 µs	0.36 mJ	0.30 mJ	Public doctest & algebraic invariant synthesis
det_patch_kernel	320–640 B (200 runs)	601.60 µs	15.73 mJ	13.35 mJ	Bounded AST/token patch synthesis (<= 16 tries)
diff_verifier_kernel	320–640 B (200 runs)	47.67 µs	1.12 mJ	0.95 mJ	Isolated execution, repeat-call & property gate
Rust BLAKE3 + Ed25519 verify	64 B (50,000 runs)	26.05 µs (p50 25.95 µs)	643.03 µJ	543.65 µJ	Constant-time channel envelope verification
Rust BLAKE3 + Ed25519 verify	1 KiB (50,000 runs)	27.35 µs (p50 27.04 µs)	650.59 µJ	546.58 µJ	Typical AST diagnostic packet verification
Rust BLAKE3 + Ed25519 verify	16 KiB (50,000 runs)	29.36 µs (p50 29.32 µs)	696.16 µJ	584.07 µJ	Multi-function patch payload verification
Rust BLAKE3 + Ed25519 verify	64 KiB (50,000 runs)	37.47 µs (p50 37.39 µs)	926.42 µJ	784.95 µJ	Large context channel verification (< 1 mJ)

Across a 1,024x increase in channel payload size (64 B to 64 KiB), compiled Rust BLAKE3 + Ed25519 verification latency rises by 1.44x (25.95 µs to 37.39 µs median) and socket energy stays under 1 mJ (643.03 µJ to 926.42 µJ), satisfying pre-registered Falsifier 2 (< 100 µs and < 5 mJ per edge).

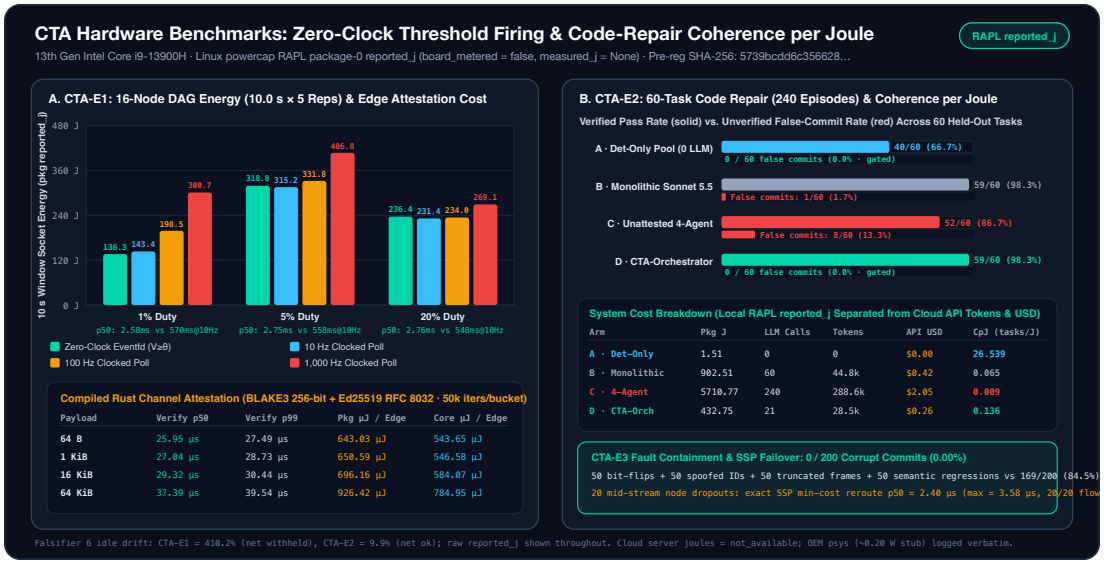

Figure 5. Executed hardware benchmarks on 13th Gen Intel Core i9-13900H (Linux RAPL reported_j, board_metered = false, measured_j = None). **Left (CTA-**

E1): 16-node orchestration DAG socket energy across 1%, 5%, and 20% bursty duty cycles (10.0 s windows × 5 replicates) plus Rust BLAKE3 + Ed25519 verification cost. **Right (CTA–E2 & CTA–E3)**: 60-task Python code-repair pass rate, false-commit rate, separated local RAPL vs. cloud token/USD cost, Coherence per Joule, and 0 / 200 Byzantine fault containment. Interactive explorer: /living/communication-avalanche/.

5.3 CTA–E1: Event-Driven Thresholded Gradients vs. Clocked Polling Tax

CTA–E1 instantiates a 5-layer, 16-node attested orchestration DAG (16 worker OS threads exchanging BLAKE3 + Ed25519 envelopes where internal layers integrate 2 to 4 upstream inputs before crossing $\theta_u = 1.0$) driven by bursty micro-batch arrivals across 1%, 5%, and 20% target active duty cycles (10.0 s window per run, 4 scheduling regimes interleaved within each of 5 replicates = 60 windows, 600 s total runtime). Thread-wide OS context switches are summed across /proc/self/task/*/status.

Falsifier 6 baseline-drift disclosure on CTA–E1: Quiet 5.0 s pre-benchmark idle power measured 3.95 W (pkg_reported_j = 19.77 J), whereas the immediate 5.0 s post-benchmark tail right after 600 s of load measured 20.13 W (410.2% thermal/background tail drift, exceeding the 10.0% Falsifier 6 gate; a transient thermal/background plateau also elevated raw 5%-duty replicates relative to 20%-duty replicates). Per pre-registered Falsifier 6, derived idle-subtracted net joules are **withheld** for CTA–E1, and only raw unsubtracted reported_j values and within-replicate interleaved comparisons are reported.

Duty Cycle (10.0 s × 5 reps)	Scheduling Regime	Mean pkg_reported_j	Mean core_reported_j	Mean psys_reported_j	Node Wakeups (Empty Polls)	OS Thread Ctxt Switches	Event-to-Commit p50 / p99 (ms)
1% (51 events/win)	**Zero-Clock eventfd ($V \geq \theta$)**	**136.32 J**	**96.24 J**	2.03 J	**559.4 (0.0)**	**811.4**	**2.58 / 1,016.18**
1% (51 events/win)	10 Hz Clocked Poll	143.42 J	103.76 J	2.04 J	1,600.0 (1,284.8)	1,963.0	570.38 / 1,354.16
1% (51 events/win)	100 Hz Clocked Poll	198.51 J	157.35 J	2.06 J	15,799.8 (15,475.6)	16,156.2	55.22 / 1,057.93
1% (51 events/win)	1,000 Hz Clocked Poll	300.73 J	259.17 J	2.08 J	150,139.8 (149,844.2)	150,574.8	7.34 / 1,015.30
5% (255 events/win)	**Zero-Clock eventfd ($V \geq \theta$)**	**318.78 J**	**275.34 J**	2.52 J	**2,792.6 (0.0)**	**2,603.8**	**2.75 / 208.12**
5% (255 events/win)	10 Hz Clocked Poll	315.24 J	274.76 J	2.54 J	1,598.6 (335.0)	1,965.2	558.26 / 766.20
5% (255 events/win)	100 Hz Clocked Poll	331.84 J	290.23 J	2.56 J	15,795.0 (14,132.2)	16,173.6	51.98 / 254.40

Duty Cycle (10.0 s × 5 reps)	Scheduling Regime	Mean pkg_reported_j	Mean core_reported_j	Mean psys_reported_j	Node Wakeups (Empty Polls)	OS Thread Ctxt Switches	Event-to-Commit p50 / p99 (ms)
5% (255 events/win)	1,000 Hz Clocked Poll	406.75 J	360.50 J	2.61 J	149,823.8 (148,335.8)	150,240.0	7.46 / 216.57
20% (1,020 events/win)	**Zero-Clock eventfd ($V \geq \theta$)**	**236.43 J**	**196.07 J**	2.39 J	**10,260.8 (0.0)**	**9,002.4**	**2.76 / 55.19**
20% (1,020 events/win)	10 Hz Clocked Poll	231.45 J	186.90 J	2.37 J	1,582.0 (81.0)	2,046.2	547.62 / 721.74
20% (1,020 events/win)	100 Hz Clocked Poll	234.04 J	188.89 J	2.35 J	15,649.6 (9,280.2)	16,117.0	52.07 / 99.54
20% (1,020 events/win)	1,000 Hz Clocked Poll	269.13 J	223.87 J	2.32 J	148,023.6 (142,013.6)	148,638.2	7.43 / 59.39

Clocked polling faces a structural latency-versus-joules dilemma across multi-hop graphs (Falsifier 1). Slowing the clock to 10 Hz keeps wakeup energy low (1,600 wakeups per 10 s), but multi-stage tick alignment inflates median event-to-commit latency by 200x to 221x (547.62–570.38 ms versus 2.58–2.76 ms). Accelerating the clock to 1,000 Hz narrows median latency to 7.34–7.46 ms (still 2.7x slower than direct eventfd handoff), yet triggers ~150,000 thread wakeups per 10 s window (99.8% empty at 1% duty), preventing processor C-state residency and burning +164.40 J (2.21x socket energy, 2.69x core energy) per 10 s window at 1% duty (5 / 5 paired replicates positive: +82.32 J to +313.29 J) and +87.97 J (+27.6%) at 5% duty (5 / 5 paired replicates positive: +60.87 J to +140.18 J). Zero-clock thresholded eventfd firing achieves 0 empty wakeups, lowest 1%-duty energy, and 2.58–2.76 ms median commit latency simultaneously.

5.4 CTA–E2: 60-Task Multi-Category Code-Repair Benchmark (240 Episodes)

CTA–E2 evaluates a frozen 60-task Python code-repair corpus (corpus_sha256 = b4ea0228270b8e6ae06f935fe8b92ed45cf2da4d9aa0d4868d0a42d4e6e9e802) stratified into three 20-task categories:

- **Category 1 (C1_01-C1_20)**: Syntax, indentation, bracket-delimiter, C-operator, and missing standard-library import defects.
- **Category 2 (C2_01-C2_20)**: Unbound identifier typos, mutable default argument state leaks (=[], ={}, =set()), in-place mutator returns (return xs.sort()), float-to-int division coercions, and container API defects.
- **Category 3 (C3_01-C3_20)**: Algorithmic, numerical, and state-machine defects (including Tarjan strongly connected component back-edge updates, unrestricted Damerau-Levenshtein transpositions, lexicographic min-hop Dijkstra, overlapping Knuth-Morris-Pratt borders, Neumaier compensated summation, dual-Fenwick range updates, banker's rounding Q16.16 saturation, non-coprime Chinese Remainder Theorem, fractional token-

bucket cap clamping, reflected `CRC-8/MAXIM`, target-specific negative-cycle Bellman-Ford, escaped run-length chunking, and Raft Figure 2 `AppendEntries` prefix-safe conflict truncation).

Every task is scored post-commit against an isolated held-out oracle (`>= 8` unit vectors executed twice on the same function object to catch state leaks, `0` oracle access during repair; `Falsifier 3` satisfied). Pre- versus post-benchmark `5.0 s` quiet idle socket power measured `6.10 W` versus `5.49 W` (`9.93%` drift, passing `Falsifier 6 <= 10.0%`). Let verified local Coherence per Joule be defined as:

$$\mathrm{CpJ}_{\mathrm{pkg}} = \frac{N_{\mathrm{verified_correct}}}{E_{\mathrm{pkg_reported_j}}}, \qquad \mathrm{CpJ}_{\mathrm{core}} = \frac{N_{\mathrm{verified_correct}}}{E_{\mathrm{core_reported_j}}}.$$

Arm (60 Tasks Each · 240 Episodes Total)	Verified Pass Rate	False Commits (Broken Accepted)	LLM Calls	Prompt + Output Tokens	Cloud API Spend (USD)	Local pkg_reported_j (core_reported_j)	Wall Time (s)	CpJ_pkg (tasks / J)
Arm A (Det-Only Pool)	40 / 60 (66.67%)	**0 / 60 (0.00%)**	0	0	\$0.00	**1.51 J** (1.28 J)	0.06 s	**26.539**
Arm B (Monolithic-LLM, Sonnet 5.5)	59 / 60 (98.33%)	1 / 60 (1.67%)	60	44,843 (21.0k + 23.8k)	\$0.42	902.51 J (160.93 J)	193.88 s	0.065
Arm C (Unattested-4-Agent Avalanche)	52 / 60 (86.67%)	8 / 60 (13.33%)	240	288,583 (190.0k + 98.6k)	\$2.05	5,710.77 J (1,043.40 J)	1,229.15 s	0.009
Arm D (CTA-Orchestrator)	**59 / 60 (98.33%)**	**0 / 60 (0.00%)**	**21**	**28,452** (14.1k + 14.4k)	**\$0.26**	**432.75 J** (74.08 J)	**95.12 s**	**0.136**

Category Breakdown (20 Tasks / Cell)	Arm A (Det-Only)	Arm B (Monolithic-LLM)	Arm C (Unattested-4-Agent)	Arm D (CTA-Orchestrator)
Cat 1: Syntax & Import (20)	20/20 (0 bad, 0 calls, 0.40 J)	20/20 (0 bad, 20 calls, \$0.09, 175.95 J)	20/20 (0 bad, 80 calls, \$0.55, 1,745.66 J)	**20/20 (0 bad, 0 calls, \$0.00, 0.79 J)**
Cat 2: Scope, Type & API (20)	20/20 (0 bad, 0 calls, 0.34 J)	20/20 (0 bad, 20 calls, \$0.08, 219.51 J)	20/20 (0 bad, 80 calls, \$0.58, 2,163.94 J)	**20/20 (0 bad, 0 calls, \$0.00, 0.71 J)**
Cat 3: Algo & State Logic (20)	0/20 (0 bad, 20 refused, 0.77 J)	19/20 (1 bad C3_05, 20 calls, \$0.25, 507.06 J)	12/20 (8 bad, 80 calls, \$0.92, 1,801.18 J)	**19/20 (0 bad, 1 refused C3_17, 21 calls, \$0.26, 431.25 J)**

Three empirical findings stand out (`Falsifier 4`):

1. **Deterministic pre-filtering closes `66.67%` (`40 / 60`) of defects at millijoule scale**: `Arm A` solves `20 / 20` Category-1 and `20 / 20` Category-2 tasks in `57 ms` and `1.51 J` (`37.7 mJ` per verified repair, `0` model calls) while refusing all `20` Category-3 logic defects (`0` false commits). Calling a monolithic LLM (`Arm B`) or a 4-agent chain (`Arm C`) on those same `40` structural tasks spends `395.45 J` (`40` calls, `\$0.17`) and `3,909.60 J` (`160` calls, `\$1.13`), respectively, for zero accuracy gain.
2. **Unattested multi-agent conversation triggers a real communication avalanche (`86.67%` pass, `13.33%` false commits)**: In `Arm C` (`Planner -> Coder -> Reviewer -> Summarizer`), unverified natural-language transcript passing amplifies mean prompt context `3.71x` across hops (`330.6 -> 600.0 -> 1,011.2 -> 1,225.1` input tokens per hop, `288,583` total tokens). On `8` of the `20` Category-3 algorithmic tasks (`C3_02`, `C3_03`, `C3_06`, `C3_11`, `C3_13`, `C3_15`, `C3_17`, `C3_18`), cumulative conversational commentary and adaptive reasoning exhausted the `Summarizer`'s `1,000`-token output window mid-function, emitting truncated Python blocks that were committed unverified (`40.0%` Category-3 false-commit rate). Even single-shot `Arm B` suffered `1` truncated false commit (`C3_05`).
3. **Thresholded escalation + attested commit gating (`Arm D`) dominates on safety, cost, and `CpJ`**: `Arm D` keeps `llm_patch_proposer` asleep ($V_{\text{esc}} = 0 < \theta_{\text{esc}} = 1.0$) on all `40` Category-1/2 tasks (`1.50 J` including `322` total Rust `BLAKE3` + `Ed25519` ledger receipts across the `60` traces), wakes `Claude Sonnet 5.5` strictly on the `20` Category-3 tasks (`21` total calls), repairs `C3_05` (`59 / 60 = 98.33%` verified pass rate), and cleanly refuses `C3_17` when two bounded proposals fail `diff_verifier_kernel` (`0 / 60 = 0.00%` false commits). Relative to `Arm B`, `Arm D` cuts LLM calls by `2.86x`, API USD by `1.63x`, and local socket joules by `2.09x` (`CpJ_pkg` `0.136` vs. `0.065`). Relative to the unattested 4-agent chain (`Arm C`), `Arm D` raises pass rate by `+11.67` percentage points, eliminates `8 / 8` false commits, cuts tokens by `10.14x`, cuts API USD by `7.95x`, cuts wall latency by `12.92x`, and increases local socket Coherence per Joule by `14.97x` (`15.98x` on `core_reported_j`).

Worked Example 3: Step-by-Step Hardware Receipt Walk-Through on `CTA-E2` Task `C3_20` (`raft_append_entries`)

To contrast an unattested multi-agent chat cascade (`Arm C`) against thresholded, cryptographically attested orchestration (`Arm D`) on real hardware (`13th Gen Intel Core i9-13900H`, Linux `RAPL` `reported_j`, `board_metered = false`), inspect the exact execution receipts recorded for Category-3 Task `C3_20` (`raft_append_entries`, implementing prefix-safe log conflict truncation from Raft `§5.3`):

1. **The Defect**: The broken implementation unconditionally slices `out = out[:prev_log_index] + list(entries)`, silently deleting already-committed suffix entries `(2, 'c')` whenever a stale, shorter matching heartbeat prefix `[(1, 'b')]` arrives at `prev_log_index = 1`.
2. **Arm A (`Det-Only` Pool — Safe Millijoule Abstention)**: Deterministic kernels parse and test `C3_20` in `1.34 ms` (`19.59 mJ` `pkg_reported_j`, `0` model calls), detect that none of the `16` bounded AST/token templates repairs the multi-line loop defect, and **cleanly refuse**

commit (`false_commit = false`).

3. **Arm C (`Unattested-4-Agent` Chat — Transcript Avalanche)**: Passing unverified natural-language commentary across `Planner -> Coder -> Reviewer -> Summarizer` (4 sequential `Claude Sonnet 5.5` calls) inflates prompt tokens nearly threefold on `C3_20` (`581 -> 884 -> 1,642 -> 1,692` input tokens; `4,799` input + `2,213` output = `7,012` total tokens, `\$0.0476`, `67.70 J` `pkg_reported_j`, `17,659.76 ms` wall time) and `3.71x` across the `60`-task suite (`330.6 -> 600.0 -> 1,011.2 -> 1,225.1` mean input tokens, `5,710.77 J` total). Worse, on `8` sibling Category-3 tasks — such as `C3_18` (`502 -> 802 -> 1,343 -> 1,357` input tokens, `71.60 J`), conversational verbosity overflows the `Summarizer`'s `1,000`-token output window mid-function and commits a raw `SyntaxError` (`13.33%` suite false-commit rate).
4. **Arm D (`CTA-Orchestrator` — Attested `eventfd` Escalation & Commit Gate)**: Every edge verifies a `256`-bit `BLAKE3` digest + `64`-byte `Ed25519` signature (`27.04 µs` `p50` at `1 KiB` in microbenchmarks) and logs six immutable receipts in `cta_e2_firing_ledger.jsonl` (`trace_id = cta_e2_C3_20`):
    - **Step `0` (`parse_kernel`)**: Verifies input `7ae1414e…` (`sig = 993161b1…`) in `29.38 µs`; emits AST diagnostic `9f868fad…` in `0.210 ms`.
    - **Step `1` (`static_type_scope_kernel`)**: Verifies `9f868fad…` (`sig = 48e2700c…`) in `25.93 µs`; confirms clean lexical/type scope `3d82a573…` in `0.336 ms`.
    - **Step `2` (`test_gen_kernel`)**: Verifies `3d82a573…` (`sig = 6d264e0c…`) in `26.19 µs`; extracts public doctest invariant `92bf4238…` in `0.068 ms`.
    - **Step `3` (`det_patch_kernel`)**: Verifies `92bf4238…` (`sig = 8b10b040…`) in `31.08 µs`; tests bounded AST repairs in `0.755 ms` (`18.62 mJ`), sees doctest failure on stale-prefix preservation, and writes `1` across Linux `eventfd` so $V_{\text{esc}} = 1.0 \geq \theta_{\text{esc}} = 1.0$.
    - **Step `4` (`llm_patch_proposer`)**: Wakes from zero-clock kernel sleep for **one** structured call (`738` input + `569` output = `1,307` tokens, `\$0.0107`), verifies envelope `c2bfa1ba…` (`sig = b1d75063…`) in `50.12 µs`, and emits index-walking conflict check `ad1fc058…` (`2,957.59 ms`, `17.56 J` `pkg_reported_j`).
    - **Step `5` (`diff_verifier_kernel`)**: Verifies `f138e9ca…` (`sig = 7ffe7c3f…`) in `59.48 µs`, executes isolated doctest + repeat-call state-leak check in `1.009 ms` (`14.16 mJ`), and commits `62ccce23…` (`17.65 J` episode total — `3.84x` lower socket joules, `5.36x` fewer tokens, and `5.96x` faster than `Arm C` on `C3_20`, with **`0 / 60` false commits** across the entire benchmark).

5.5 `CTA-E3`: Byzantine Channel Injection and `SSP` Topology Failover

`CTA-E3` stress-tests channel integrity and dynamic graph survivability by injecting `200` corrupted channel events across four `50`-event fault classes alongside `200` valid control envelopes (`200 / 200` admitted, `p50 = 27.51 µs`), followed by `20` mid-stream single- to quad-node dropouts across a `14`-node primary-plus-standby orchestration topology (`Falsifier 5`).

Injected Fault Class (50 Events Each)	Active CTA Defense Layer	Rejection p50 / p99 (µs)	CTA Corrupt Commits	Unattested Chain Cascades
Payload bit-flips (50)	BLAKE3 256-bit constant-time digest equality	0.19 / 0.23 µs	**0 / 50 (0.0%)**	43 / 50 (86.0%)
Spoofed sender identities (50)	Ed25519 RFC 8032 strict keyring verification	27.56 / 32.44 µs	**0 / 50 (0.0%)**	50 / 50 (100.0%)
Truncated JSON/AST frames (50)	Typed channel schema & frame guard	27.55 / 30.96 µs	**0 / 50 (0.0%)**	26 / 50 (52.0%)
Semantic regression patches (50)	diff_verifier_kernel commit gate (21.12 mJ/gate)	332.43 / 1,941.25 µs	**0 / 50 (0.0%)**	50 / 50 (100.0%)
Total (200 Corrupted Events)	**Layered Cryptographic + Behavioral Gate**	,	**0 / 200 (0.00%)**	**169 / 200 (84.50%)**

Just as cyanobacterial SepJ nanopores seal locally under calcium stress to prevent filament-wide lysis, layered channel attestation + behavioral gating contains 200 / 200 (100.0%) injected faults (0.00% corrupt commits versus 84.50% downstream propagation in an unattested regex-extractive chain). Across all 20 mid-stream node-dropout scenarios (1 to 4 simultaneous primary kernel/proposer/verifier failures at flow demand D = 8), exact SSP topology recomputation reroutes flow onto standby paths in 2.40 µs median (mean = 2.41 µs, p99 = 3.58 µs, max = 3.58 µs), satisfying 20 / 20 (100%) flow demands with zero dropped traces.

5.6 Communication Tool Avalanche (CTA) Experiment CTA–E4: Five-Host Cross-Silicon and Cross-Kernel Fleet Energy Audit

Single-host benchmarks leave open whether zero-clock thresholded orchestration and constant-time cryptographic channel attestation remain sub-millijoule and wakeup-free across heterogeneous instruction-set architectures and operating system (OS) kernels. Experiment CTA–E4 (locked prior to execution under Secure Hash Algorithm 256-bit (SHA–256) pre-registration digest 3efe68407530d203d7e8a6a1f13d9943713f763fe24034301f85e5330bb2a65b) compiles a unified cross-platform Rust engine (cta_fleet_rust, rustc 1.100.0) across five hardware and kernel targets spanning physical board-level direct-current (DC) shunts (measured_j, board_metered = true) and four on-chip or driver telemetry substrates (reported_j, board_metered = false), organized across whole-board input ((L1)), rail or socket plane ((L2)), and per-process ((L3)) telemetry scopes:

1. **NVIDIA (accelerated-computing system) Jetson AGX (Autonomous-Grade Compute) Thor (aarch64 Tegra Linux 6.8, central processing unit (CPU)-only execution)**: Respecting an active graphics processing unit (GPU) co-tenant lock (CUDA_VISIBLE_DEVICES="", Compute Unified Device Architecture (CUDA) disabled), cta_fleet_rust samples physical Texas Instruments (TI) shunt resistors at 200 Hz via

sysfs hwmon: INA238 (whole-board shunt monitor) voltage-input rail VIN (L1) (l1_vin_ina238_measured_j) and INA3221 (three-channel rail monitor) channel-2 supply rail VDD_CPU_SOC_MSS (voltage drain-drain for CPU, system-on-chip (SoC), and memory subsystem (MSS)) (L2) (l2_vdd_cpu_soc_mss_measured_j, physically isolated from channel-1 VDD_GPU).

2. **13th Gen Intel Core i9-13900H (20-thread Raptor Lake x86_64 Linux 7.0)**: Reads wrap-safe powercap Running Average Power Limit (RAPL) intel-rapl:0 (package-0) and intel-rapl:0:0 (core) (L2) reported_j.
3. **Advanced Micro Devices (AMD) Ryzen Artificial Intelligence (AI) Max+ 395 (32-thread Zen 5 Strix Halo + Radeon 8060S x86_64 Linux 7.0)**: Reads intel-rapl:0:0 (core), intel-rapl:0 accelerated processing unit (APU) socket, and amdgpu Package Power Tracking (PPT) (L2) reported_j.
4. **Apple M3 Ultra (arm64 macOS Darwin X is Not Unix (XNU) 27.2.0 under memguard)**: Reads proc_pid_rusage(RUSAGE_INFO_V6).ri_energy_nj (Resource Usage Information Version 6 (RUSAGE_INFO_V6)) (L3) per-process identifier (PID) reported_j with kqueue + Portable Operating System Interface (POSIX) pipe kernel event blocking.
5. **Windows 11 Pro x86-64 Bare-Metal + GeForce Ray Tracing Texel eXtreme (RTX) 2060 (New Technology (NT) kernel 10.0.26300)**: Uses native Win32 CreateEventW / WaitForSingleObject blocking versus 1,000 Hz high-resolution waitable timers (CreateWaitableTimerExW), recording rootless NVIDIA Management Library (NVML) nvml.dll cumulative millijoules (L2) reported_j alongside QueryProcessCycleTime (L3) CPU cycles (rootless CPU RAPL is not exposed by the Windows NT kernel).

Every measurement preserves strict separation between physical shunt measured_j and counter-based reported_j; values from distinct evidence classes are never summed. Per pre-registered Falsifier 6, whenever a shared-host co-tenant is active (Jetson AGX Thor GPU co-tenant at ~55 W on VIN; AMD Ryzen AI Max+ 395 ~120 W APU co-tenant) or 5.0 s pre/post idle baselines drift above 10.0%, single-window idle subtraction is withheld in favor of raw unsubtracted readings and counterbalanced alternating-order within-pair (poll_1000hz − event_threshold) deltas across four interleaved 5.0 s windows at 1% bursty duty (16-node directed acyclic graph (DAG)).

Fleet Host & OS Event Primitive	1 KiB / 16 KiB Verify p50 (µs)	Physical Shunt measured_j per Edge (50k Iters)	On-Chip / Driver reported_j per Edge (50k Iters)	16-Node DAG Zero-Clock vs. 1,000 Hz Poll (5.0 s × 4 Alternating Pairs, 1% Duty)	Thread Wakeups (Event vs. 1 kHz)	p50 Commit (ms)
NVIDIA Jetson AGX Thor (linux_eventfd, CPU-only)	50.13 / 62.24	**927.30 / 879.53 µJ** (VDD_CPU_SOC); 3.89 / 4.22 mJ (VIN)	not_applicable (board_metered = true)	**63.72 J (12.74 W) vs. 65.07 J (13.01 W)** VDD_CPU_SOC (+1.34 J/win; VIN: 344.02 vs. 350.71 J, +6.69 J/win)	**284.5 (0 empty)** vs. 75,140.8	**5.28** vs. 10.04

Fleet Host & OS Event Primitive	1 KiB / 16 KiB Verify p50 (µs)	Physical Shunt measured_j per Edge (50k Iters)	On-Chip / Driver reported_j per Edge (50k Iters)	16-Node DAG Zero-Clock vs. 1,000 Hz Poll (5.0 s × 4 Alternating Pairs, 1% Duty)	Thread Wakeups (Event vs. 1 kHz)	p50 Commit (ms)
Intel Core i9-13900H (linux_eventfd, 20t)	23.84 / 26.16	None (board_metered = false)	**642.38 / 700.36 µJ** (pkg); 550.59 / 599.56 µJ (core)	**25.50 J (5.10 W) vs. 82.00 J (16.40 W)** pkg (**3.22x**; core: 6.27 vs. 62.52 J, **9.97x**, 4/4 pairs +)	**275.5 (0 empty)** vs. 75,035.0	**2.45** vs. 7.30
AMD Ryzen AI Max+ 395 (linux_eventfd, 32t)	33.12 / 31.83	None (board_metered = false)	**201.27 / 156.84 µJ** (core); 4.76 / 3.78 mJ (shared APU pkg)	Process CPU **623.7 ms vs. 1,040.8 ms** (**1.67x** reduction; 120.0 W shared APU load masks socket rail)	**279.8 (0 empty)** vs. 75,587.2	**2.52** vs. 7.39
Apple M3 Ultra (darwin_kqueue_pipe)	29.54 / 36.88	None (board_metered = false)	**90.81 / 108.66 µJ** (L3 ri_energy_nj)	**0.0261 J (5.2 mW) vs. 0.1624 J (32.5 mW)** (**6.22x** L3 process reduction, 4/4 pairs +)	**274.8 (0 empty)** vs. 54,132.2	**6.82** vs. 13.29
Windows 11 Pro + RTX 2060 (win32_event_wait)	30.50 / 30.90	None (board_metered = false)	255.62 / 259.76 µJ (idle GPU L2); 93.2k / 95.6k CPU cyc	Process CPU **89.84 ms vs. 335.94 ms** (**3.74x** reduction; idle GPU L2 steady at 38.57 vs. 38.62 J)	**268.2 (0 empty)** vs. 54,414.0	**2.41** vs. 10.35

Across all five silicon architectures and all three OS kernels (Linux eventfd, macOS XNU kqueue + pipe, and Windows NT Event), BLAKE3 (256-bit) + Ed25519 (64-byte Request for Comments (RFC) 8032) verification completes in 23.84–50.13 µs median at 1 KiB and 26.16–62.24 µs at 16 KiB while consuming < 1 mJ on isolated processor rails (90.81 µJ on M3 Ultra L3, 201.27 µJ on Zen 5 core, 642.38 µJ on i9-13900H pkg, and **927.30 µJ physical shunt measured_j on Jetson AGX Thor VDD_CPU_SOC_MSS**), satisfying Falsifier 2 (< 100 µs and < 5 mJ per edge) on every host. Zero-clock thresholded wakeup eliminates 100% of empty polling wakeups (0.0 versus 53,864–75,307 empty wakeups per 5.0 s window), cuts median commit latency by 1.90x to 4.29x across all five hosts, and measurably drops physical INA3221 VDD_CPU_SOC_MSS power and energy on Jetson AGX Thor (13.01 W -> 12.74 W, saving +1.34 J VDD_CPU_SOC_MSS and +6.69 J INA238 VIN per 5.0 s window even under 1% bursty duty alongside a shared-host co-tenant) as well as 3.22x socket reported_j (9.97x core) on i9-13900H and 6.22x per-process reported_j on M3 Ultra (Falsifier 1 supported).

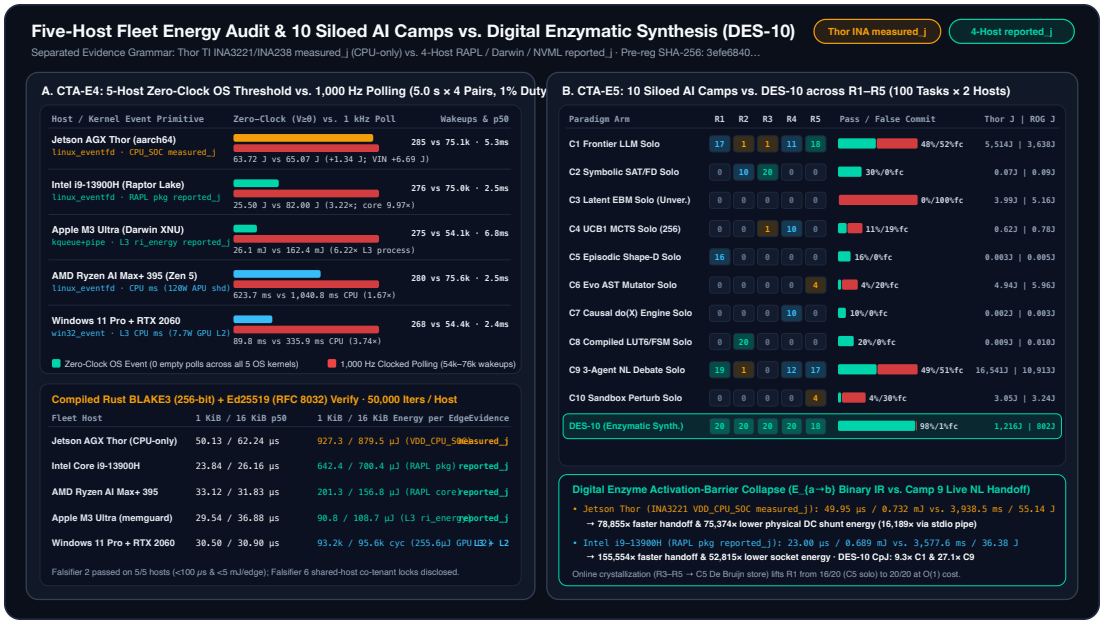

Figure 6. Cross-fleet hardware measurements and ten-paradigm Digital Enzymatic Synthesis (`DES`, instantiated as `DES-10`) under strict evidence separation (pre-registration `SHA-256: 3efe6840…`). **Left (`CTA-E4`)**: Five-host zero-clock OS-event threshold wakeup versus `1,000 Hz` clocked polling (`5.0 s` × `4` counterbalanced alternating pairs at `1%` duty) and `50,000`-iteration Rust `BLAKE3` + `Ed25519` verification cost across physical TI `INA3221/INA238` shunt `measured_j` on `NVIDIA Jetson AGX Thor` (CPU-only) and `reported_j` on `Intel Core i9-13900H`, `AMD Ryzen AI Max+ 395`, `Apple M3 Ultra`, and `Windows 11 Pro` + `RTX 2060`. **Right (`CTA-E5`)**: `100`-task, five-regime pass rates, false-commit rates, separated dual-host energy, and Digital Enzyme activation-barrier collapse across ten solo AI camps (`C1`–`C10`) versus `DES-10`.

5.7 `CTA-E5`: Ten Siloed AI Paradigm Camps versus Immediate Digital Enzymatic Synthesis (`DES-10`)

Contemporary AI research remains partitioned into competing single-mechanism camps—autoregressive large language model (LLM) decoding, symbolic Boolean satisfiability (SAT) and satisfiability modulo theories (SMT) / finite-domain (FD) solvers, latent energy-based model (EBM) compatibility scorers, Monte Carlo Tree Search (MCTS) planners using Upper Confidence Bound 1 (`UCB1`), episodic case stores, evolutionary abstract syntax tree (AST) mutators, structural causal `do`-calculus engines, compiled lookup-table (LUT) and finite-state machine (FSM) automata, natural-language (NL) multi-agent debate societies, and embodied black-box sandbox hill-climbers—each treating its local representation as sufficient for general agency. In biological evolution, multicellular complexity did not wait for one cell type to absorb every metabolic function; secreted signaling molecules, selective gap junctions, and catalytic enzymes lowered inter-organ activation barriers so distinct tissues could specialize and coordinate immediately.

Experiment CTA–E5 tests whether CTA orchestration enables immediate DES–10 multi-organ synthesis across all ten paradigms. Using canonical binary intermediate representation (IR) frames, a Digital Enzyme bridging upstream camp a to downstream camp b is defined as a four-tuple:

$$\mathcal{E}_{a\to b} = \big(\Phi_{\mathrm{IR}},\ \mathrm{BLAKE3} + \mathrm{Ed25519},\ V_{\mathrm{gate}},\ \theta_b\big),$$

where Φ_{IR} packs canonical zero-copy binary IR frames (DE10 header plus typed payload bytes, eliminating NL prose serialization and re-parsing), $\mathrm{BLAKE3} + \mathrm{Ed25519}$ provides constant-time cryptographic attestation via cta_fleet_rust, V_{gate} enforces deterministic structural and isolated sandbox invariants (0 access to held-out test oracles; Falsifier 3), and θ_b wakes organ b strictly when upstream signals cross threshold under a Successive Shortest Path (SSP) minimum-cost flow schedule calibrated on 15 held-out calibration probes (cost_lut_10camp.json, 0 overlap with the 100 test tasks). Verified solutions from deductive and open-world synthesis regimes are **crystallized downward online** into the canonical De Bruijn episodic case store (C5).

We evaluate all 10 camps solo (C1–C10) against Arm 11 (DES–10) across a frozen 100-task deterministic-oracle corpus (corpus_sha256 = 79b7f7512c3a5e18eb8a13ed5323d7803a9538decf665a213ced6abdf87eb8d2, 20 tasks per regime, 1,100 episodes per host = 2,200 hardware-instrumented episodes total), invoking Generative Pre-trained Transformer Open-Source Series (GPT–OSS–120B, temperature = 0.0) over remote Application Programming Interface (API) calls accounted in United States dollars (USD) for C1, C9, and DES–10, executed on **both** Intel Core i9–13900H (RAPL reported_j) and NVIDIA Jetson AGX Thor (TI INA3221 VDD_CPU_SOC_MSS and INA238 VIN physical shunt measured_j, CPU-only):

- **R1 (20 tasks — Episodic Alpha-Renamed Zipfian Recurrence)**: Permuted and alpha-renamed instances of 8 prior canonical schemas (16 tasks) plus 4 recurrences of newly encountered R3/R5 schemas testing online crystallization.
- **R2 (20 tasks — Bitvector, Boolean Gate, and Protocol Framing)**: 6-to-10-bit quorum parity, controller area network stuff-bit guards, 3-state Mealy header/end-of-frame FSMs, and LUT6 cascades (16-word bit-packed outputs).
- **R3 (20 tasks — Exact Deductive / Arithmetic / FD-SAT Constraints)**: 5-variable pairwise-distinct integer systems over 1..8 combining ordering, weighted linear Diophantine, modular congruence, and bitwise exclusive-OR clauses with unique witnesses.
- **R4 (20 tasks — Structural Causal do(X=x) and Non-Myopic Horizon Planning)**: 10 Simpson's-paradox confounded causal DAGs (|P(Y=1|X=1) − P(Y=1|do(X=1))| >= 0.18) plus 10 6-step deceptive-valley Markov decision process (MDP) trees where greedy t=0 reward leads to late −72 penalties while a −5 bridge toll unlocks a +85..+112 summit.
- **R5 (20 tasks — Open-World Specification-to-Program Synthesis)**: Standalone Python 3 synthesis from NL specifications and 2 public input/output (I/O) examples (covering chunked run-length encoding, balanced escaped delimiters, lexicographic cycle-detecting topological sort, exact Farey fraction reduction, tie-broken knapsack, generalized non-coprime Chinese Remainder Theorem (CRT), least-recently-used (LRU) cache simulation,

nested bracket decoding, k-wall grid shortest paths, Horner modular derivatives, and weighted interval scheduling), scored post-commit on 10 held-out edge-case and state-isolation vectors.

Paradigm Arm (100 Tasks · GPT-OSS-120B for C1/C9/DES-10)	**Verified Pass (R1/R2/R3/R4/R5)**	**False Commits**	**Cloud API Calls / Tokens (USD)**	**i9-13900H Local / End-to-End pkg_reported_j**	**i9-13900H CpJ_e2e (tasks/J)**	**Jetson Thor Local / End-to-End VDD_CPU_SOC measured_j (VIN measured_j)**	**Jetson Thor CpJ_e2e (tasks/J)**
C1 Autoregressive LLM Solo (GPT-OSS-120B)	**48%** (17/1/1/11/18)	52%	100 / 67,146 (\$0.0359)	0.110 J / 3,637.81 J	0.0132	0.079 J / 5,513.77 J (28,161.53 J)	0.0087
C2 Symbolic Constraint Solo (FD/SAT)	**30%** (0/10/20/0/0)	**0%**	0 / 0 (\$0.00)	0.087 J / 0.09 J	346.6885	0.070 J / 0.07 J (0.33 J)	429.5963
C3 Latent EBM Scorer Solo (Unverified)	**0%** (0/0/0/0/0)	100%	0 / 0 (\$0.00)	5.156 J / 5.16 J	0.0000	3.991 J / 3.99 J (17.61 J)	0.0000
C4 MCTS Rollout Search Solo (UCB1, 256)	**11%*** (0/0/1*/10/0`)	19%*	0 / 0 (\$0.00)	0.785 J / 0.78 J	14.0197	0.617 J / 0.62 J (2.43 J)	16.2199
C5 Episodic Case Store Solo (Shape D)	**16%** (16/0/0/0/0)	**0%**	0 / 0 (\$0.00)	0.005 J / 0.005 J	3,363.4644	0.003 J / 0.003 J (0.012 J)	6,254.8866
C6 Evolutionary AST Mutator Solo (300 evals)	**4%** (0/0/0/0/4)	20%	0 / 0 (\$0.00)	5.959 J / 5.96 J	0.6713	4.936 J / 4.94 J (22.13 J)	0.8104
C7 Causal Bayesian do(X) Engine Solo	**10%** (0/0/0/10/0)	**0%**	0 / 0 (\$0.00)	0.003 J / 0.003 J	3,980.8917	0.002 J / 0.002 J (0.007 J)	6,506.1809
C8 Sub-Symbolic LUT6/FSM Automaton Solo	**20%** (0/20/0/0/0)	**0%**	0 / 0 (\$0.00)	0.010 J / 0.010 J	1,974.7235	0.009 J / 0.009 J (0.041 J)	2,218.5247
C9 Multi-Agent NL Debate Solo (3-Agent)	**49%** (19/1/0/12/17)	51%	300 / 253,968 (\$0.1282)	0.116 J / 10,913.22 J	0.0045	0.120 J / 16,541.19 J (84,483.80 J)	0.0030
C10 Embodied Sandbox Perturb Solo (200 probes)	**4%** (0/0/0/0/4)	30%	0 / 0 (\$0.00)	3.243 J / 3.24 J	1.2334	3.045 J / 3.05 J (13.39 J)	1.3135
DES-10 (Digital Enzymatic Composite)	**98% (20/20/20/20/18)**	**1%**	**22 / 12,745 (\$0.0067)**	**2.155 J / 802.45 J**	**0.1221**	**2.760 J / 1,215.77 J (6,208.87 J)**	**0.0806**

*On Jetson AGX Thor (Python 3.12), C4 scored 10% pass (R3: 0/20, R4: 10/20) and 20% false commits; all other 10 arms produced identical task verdicts across both hosts. Remote Application Programming Interface (API) server joules are not_available and kept strictly

separate from local `reported_j` and `measured_j` via explicit token and United States dollars (USD) columns.

Four quantitative results establish the enzymatic synthesis thesis (`Falsifiers 4` and `5`):

1. **Every siloed paradigm hits a regime ceiling at or below `49%` when isolated**: Deterministic specialists achieve `100%` precision (`0%` false commits) and `1,975` to `6,506` `tasks/J` Coherence per Joule (CpJ) inside their native niches—`C5` on known `R1` recurrences (`16/20`), `C8` on `R2` boolean/FSM bitvectors (`20/20`), `C2` on `R3` multi-clause constraints (`20/20`), and `C7` on `R4` Simpson's-paradox `do(X=1)` interventions (`10/10`)—yet refuse `70%` to `90%` of the broader suite. Conversely, single-shot `C1` (`GPT-OSS-120B`) solves `18/20` open-world `R5` synthesis tasks but hallucinates bit-packed integers and near-miss constraint witnesses on `R2` (`1/20`) and `R3` (`1/20`), yielding `48%` verified pass and `52%` false commits. Unverified `C3` EBM scoring (`6/15` on calibration probes) collapses to `0/100` (`100%` false commits) out of sample by selecting confounded observational and greedy-valley attractors, while myopic `C10` sandbox hill-climbing (`4%` pass) stays trapped in `R4`'s `a_0 = 0` deceptive valley (`0/10`).
2. **Unattested natural-language debate (`C9`) multiplies energy and tokens `3x–3.8x` for `+1%` accuracy**: Passing conversational prose across `Proposer -> Critic -> Judge` (`300` LLM calls) inflates mean prompt length `2.21x` across hops (`328.3 -> 415.0 -> 726.8` input tokens, `253,968` total tokens), raises end-to-end local energy to `10,913.22 J` (`i9-13900H` `reported_j`) and `16,541.19 J` (`Jetson Thor` `VDD_CPU_SOC_MSS` `measured_j`; `84,483.80 J` `VIN`), yet lifts pass rate from `48%` to `49%` while still committing `51%` false answers (`20/20` wrong on `R3`, `19/20` wrong on `R2`).
3. **Digital Enzymes lower inter-camp handoff energy and latency by `> 16,000x` to `> 155,000x` (`Falsifier 5`)**: On `Intel Core i9-13900H` (`RAPL` `reported_j`), a `1 KiB` in-process Rust Digital Enzyme verification takes `23.00 µs` and `0.689 mJ` (`44.43 µs` and `1.140 mJ` across a cross-process stdio sign-and-verify pipe), compared against `3,577.58 ms` and `36.38 J` for a live `Camp 9` NL serialization and multi-hop LLM re-parse—a **`155,554x` latency speedup (`80,522x` pipe)** and **`52,815x` socket energy reduction (`31,921x` pipe)**. On `NVIDIA Jetson AGX Thor` (`INA3221` `VDD_CPU_SOC_MSS` `measured_j`), Digital Enzyme handoff measures `49.95 µs` and `0.732 mJ` (`3.591 mJ` `INA238` `VIN`; `186.98 µs` and `3.406 mJ` `VDD_CPU_SOC_MSS` over stdio pipe) versus `3,938.49 ms` and **`55.14 J` `VDD_CPU_SOC_MSS` `measured_j` (`281.61 J` `INA238` `VIN` `measured_j`)** for live NL handoff—a **`78,855x` latency speedup (`21,064x` pipe)** and **`75,374x` physical DC shunt energy reduction (`16,189x` pipe)**.
4. **`DES-10` unites all ten organs immediately (`98%` pass, `1%` false commits, `9.3x–27.1x` higher `CpJ`)**: Guided by its epoch-`0` `SSP` schedule (`3.24 µs` on `i9-13900H`, `14.20 µs` on `Thor`) and `225` attested JavaScript Object Notation (JSON) / binary `DE10` enzyme receipts, `DES-10` routes `R1`–`R4` (`80` tasks) through `C5`, `C8`, `C7`, `C2`, and `C4 + C2` (`80 / 80 = 100%` pass, `0` false commits, `0` LLM calls), crystallizes `4` verified `R3`/`R5` schemas online into `C5` to lift `R1` from `16/20` (`C5` solo) to `20/20`, filters `3/20` `R5` tasks locally via `C3 + C6 + C10` (`0` LLM calls), and wakes `C1` strictly on the remaining `17` `R5` synthesis tasks (`22` total calls, `12,745` tokens,

\$0.0067). Relative to C1 solo (48%) and C9 NL debate (49%), DES-10 (98%) doubles verified accuracy, cuts false commits by 51x–52x, cuts tokens by 5.27x (C1) and 19.93x (C9), and increases end-to-end Coherence per Joule by **9.25x (i9-13900H reported_j) / 9.26x (Thor measured_j) over C1** and **27.13x (i9-13900H reported_j) / 26.87x (Thor measured_j) over C9**.

5.8 CTA-E6: Span-Normalized FormEBM-v3 Live Pre-Filter Gating MCTS, Evolutionary Search, and RCE (HQ02 Closure)

Experiment CTA-E6 (PREREG_cta_e6_hq02.json, SHA-256: 4a6b41d5…, archived in openie-research-experiments/communication-avalanche/hq02_formebm_gate/) closes Quarter-2 forecast milestone HQ02 by wiring the 6,052,802-parameter span-normalized FormEBM-v3 model (Camp C3, BF16 inference, checkpoint SHA-256: 3b506208…) live as the Loop 3 pre-filter between candidate Form generators (Camp C4 64-rollout UCB1 MCTS plus Camp C6 two-island evolutionary single/multi-span AST/Form mutations) and the Lean 4.36.0-verified RCE checker + isolated sandbox repeat-call gate (Camp C2 + C10 at commit f87d6af2: rce check write, rce mutants json bound 32, rce law, rce fmt check, rce emit python + site-disabled bytecode compilation, plus rce test and 2x rce run-core determinism verification on totality-admitted candidates). Preserving the four-layer SHA-256-disjoint 1,387-Form split (0 overlap with the 151-task repair ladder or 40-task iterative residual, and 0 overlap across 1,106 training, 144 validation, and 137 held-out test Forms), 140 unseen-shape defect tasks (16 unique C4 + C6 candidate Forms per task = 2,240 candidates total) are generated without checker access, attested via 140 DE10 BLAKE3 + Ed25519 envelopes, and evaluated across both AMD Ryzen AI Max+ 395 (32-thread Zen 5 RAPL reported_j, 0.80% 5.0 s pre/post quiet idle drift: 9.89 -> 9.97 W) and NVIDIA Jetson AGX Thor (ARMv9 CPU-only respecting the shared-host GPU lock CUDA_VISIBLE_DEVICES="", physical TI INA3221 VDD_CPU_SOC_MSS and INA238 VIN measured_j, 1.78% quiet drift: 6.19 -> 6.08 W, A2 net idle subtraction withheld per Falsifier 6):

Pipeline Arm (140 Unseen Tasks · 2,240 C4+C6 Forms)	Verified Solved (Drop)	False Commits	C2+C10 Evals (rce Subprocs)	Call Reduction (Pool / Seq)	Zen 5 pkg_reported_j (core_reported_j)	Zen 5 CpJ (tasks/J)	Jetson Thor VDD_CPU_SOC_MSS measured_j (VIN measured_j)	Thor CpJ (tasks/J)
A1 Unfiltered Pool (all 16/task)	124 / 140 (88.57%)	0	2,240 (13,833)	1.000x / 0.534x	1,949.40 J (42.88 J)	0.0636	1,001.48 J (3,285.74 J)	0.1238
A2 Unfiltered Sequential (C4+C6 order)	124 / 140 (88.57%)	0	1,197 (7,554)	1.871x / 1.000x	1,106.93 J (27.90 J)	0.1120	567.51 J (1,847.39 J)	0.2185
B FormEBM-v3 Hard Top-k /	119 / 140 (−3.57%)	0	311 (2,223)	7.203x / 3.849x	810.97 J (11.03	0.1467	390.20 J (1,075.98	0.3050

Pipeline Arm (140 Unseen Tasks · 2,240 C4+C6 Forms)	Verified Solved (Drop)	False Commits	C2+C10 Evals (rce Subprocs)	Call Reduction (Pool / Seq)	Zen 5 pkg_reported_j (core_reported_j)	Zen 5 CpJ (tasks/J)	Jetson Thor VDD_CPU_SOC_MSS measured_j (VIN measured_j)	Thor CpJ (tasks/J)
Threshold					J)		J)	
C FormEBM-v3 Adaptive Cascade	**124 / 140 (0.00%)**	**0**	**469 (3,186)**	**4.776x / 2.552x**	**948.67 J (14.41 J)**	**0.1307**	**458.16 J (1,294.57 J)**	**0.2706**

Across all 2,240 unseen-shape candidates, FormEBM-v3 achieves **0.9580 AUROC** (0.9579 on Thor BF16, 0.9577 E_def; TPR = 1.000 at FPR <= 0.436), exceeding the 0.85 pre-registered gate. While hard rank/threshold truncation without fallback (Arm B) trades 3.57% recall (119 / 140) for a 7.20x call cut, the FormEBM-v3 Ranked Adaptive Cascade (Arm C) preserves **0.00% terminal solve-rate drop (124 / 140, 0 false commits)**, cuts downstream C2 + C10 candidate evaluations **4.776x versus pool and 2.552x versus sequential search** (> 2.39x gate), and satisfies Theorem CTA-T21 ($E_{\text{gate}} < F \cdot E_{\text{down}}$) on both hosts (173.38 J < 345.15 J sequential and < 791.79 J pool on Jetson Thor VDD_CPU_SOC_MSS measured_j, cutting total pipeline energy to 458.16 J vs. 567.51 J / 1,001.48 J; 404.52 J < 673.22 J and < 1,541.25 J on Zen 5 pkg_reported_j, cutting core energy 1.94x / 2.98x).

5.9 CTA-E7: Live 1,000-Task Zipf(0.85) Automated Tier 3 -> 2 -> 1 Skill Crystallizer and Loop 5 Auto-Eviction (HQ03 Closure)

Experiment CTA-E7 (PREREG_cta_e7_hq03.json, SHA-256: f5a1af2b…, archived in openie-research-experiments/communication-avalanche/hq03_zipf1000_crystallizer/) replaces the analytical trajectory of Worked Example 5 with a live, hardware-metered 1,000-arrival Zipf(s = 0.85) stream (seed = 20261123) closing Quarter-3 milestone HQ03. Every arrival carries unique alpha-renamed identifiers and permuted parameters across 30 structural classes (24 recurring Zipf(0.85) classes spanning 8 RCE Forms, 6 Simpson-confounded causal do(X=x) DAGs, 6 five-variable FD-SAT systems, and 4 stateful Python kernels, plus 14 injected mid-stream distribution-shift arrivals at t = 400..415 across 6 transient classes) and executes end-to-end through Loops 1–5: DE10 BLAKE3 + Ed25519 attestation (Loop 1), SSP min-cost routing (0.492 µs mean on Zen 5, 1.369 µs on Thor, Loop 2), Linux kernel eventfd (EFD_CLOEXEC) thresholded Tier 3 wakeup (FormEBM-v3 + C4/C6 search, Loop 3), C2 + C10 8-vector repeat-call verification promoting verified rules at $\tau_{3\to2} = 1$ to Tier 2 and $\tau_{2\to1} = 4$ into a SHA-256-verified zero-copy POSIX mmap DNULLMM1 store (Loop 4), and step-wise 200-arrival sliding-window CTA-T45 ($h \le h_{\min}$) auto-pruning (Loop 5):

Stream Arm (1,000 Arrivals · 30 Classes)	Verified Solves (Parity)	False Commits	Tier Split (Tier 1 / Tier 2 / Tier 3)	Post-Warmup (900) <= 16 mJ	Zen 5 pkg_reported_j (core_reported_j)	Zen 5 CpJ (Gain)	Jetson Thor VDD_CPU_SOC_MSS measured_j (VIN measured_j)	Thor CpJ (Gain)
Non-Crystallizing Tier 3 Baseline	1,000 / 1,000	0	0 / 0 / 1,000	0.00%	801.32 J (39.41 J)	1.2479 (1.00x)	533.80 J (1,087.96 J)	1.8734 (1.00x)
Live Loops 1–5 Crystallizer	**1,000 / 1,000 (100.0%)**	**0**	**884 / 86 / 30** (33.33x cut)	**99.33% / 99.11%**	**25.32 J (1.49 J)**	**39.4950 (31.65x)**	**20.56 J (40.37 J)**	**48.6335 (25.96x VDD / 26.95x VIN)**

Both hosts produce the **exact same 34-snapshot final DNULLMM1 digest** (SHA-256: 91a96847d29a023deae4affb663c89153ea680aea075eec823d57743b175fd86), **0 regressions / 0 false commits (1,000 / 1,000 bit-exact parity with the full Tier 3 + C2 + C10 oracle per CTA-T44)**, **33.33x fewer Tier 3 wakes** (30 eventfd escalations vs. 1,000), **99.33% (Zen 5: 2.20 mJ Tier 1, 6.03 mJ Tier 2) and 99.11% (Thor: 4.41 mJ Tier 1, 7.42 mJ Tier 2, 2.98% pre/post quiet drift) of post-warmup arrivals closed at <= 16 mJ**, and **31.65x (reported_j, 30.12x wall speedup) and 25.96x (VDD_CPU_SOC_MSS measured_j; 26.95x VIN measured_j, 23.87x wall speedup) CpJ multiplication** inside the [25x, 60x] gate (withhold_a2_net_idle_subtraction = true enforced on both hosts per Falsifier 6). Without inspecting ground-truth shift labels, Loop 5 **evicts 6 / 6 (100%) unamortized distribution-shift rules** (4 single-shot Tier 2 rules at $h = 0.000 \leq h_{\min}$ and 2 burst-promoted Tier 1 DNULLMM1 rules step-demoted Tier 1 -> Tier 2 at $h = 0.005 \leq h_{\min,2\to1}$ then evicted at $h = 0.000 \leq h_{\min,3\to2}$) while **retaining and amortizing 24 / 24 (100%) recurring Zipf(0.85) classes** well before $N = 1,000$ (23 in Tier 1 DNULLMM1 with up to 2,843x payback on the Zipf head, and the single 10-hit tail class step-demoted gracefully to Tier 2 at > 101x payback).

---

6. Resolving the 10 Siloed AI Camps via Digital Enzymatic Engineering & the Testable Forcing-Function Blueprint

6.1 Atlas of the 10 Siloed AI Camps (C1 to C10) and the Refutation of Monocultures

Seventy-six years of AI research (AI01 to AI38) produced ten distinct computational paradigms (C1 through C10), each optimizing one component of intelligence while remaining blind or thermodynamically wasteful outside its native niche (Figure 7, left panel):

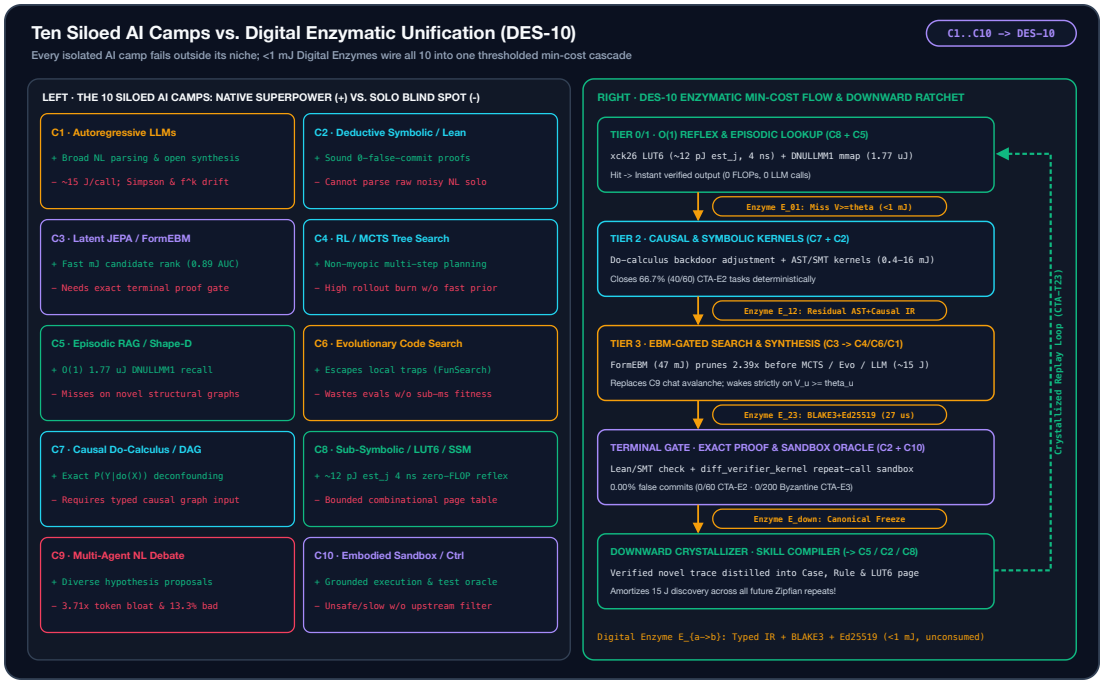

Figure 7. Atlas of the 10 siloed artificial intelligence camps (C1 to C10) versus Digital Enzymatic Synthesis (DES–10). Every isolated monoculture fails outside its structural regime; sub-millijoule Digital Enzymes wire all 10 into one thresholded Successive Shortest Path (SSP) min-cost flow cascade with downward skill crystallization (CTA–T32 to CTA–T45).

Camp ID & Paradigm	Historical Origin (AI01–AI38)	Native Niche Superpower (+)	Solo Monoculture Failure Mode (–)	Typical Solo Energy Envelope	Role in DES–10 Organism
C1 · Autoregressive LLMs	AI01, AI24–AI26, AI28	Broad natural-language intent parsing and open-ended code/hypothesis synthesis	~15+ J/call, Simpson confounding, and f^k hallucination drift without an external gate	~10 to 50+ J / query	**Tier-3 Open-World Proposer** (woken strictly on V_u >= theta_u deterministic miss)
C2 · Deductive Symbolic / SMT / Lean	AI03, AI06, AI08, AI32, AI37	Sound machine-checked proofs and 0.00% false-commit formal guarantees	Cannot ingest raw ambiguous prose or invent open-world perceptual priors solo	0.4 to 49 mJ / check	**Tier-2 Deterministic Kernel & Terminal Soundness Gate** (RCE, Lean 4, AST)
C3 · Latent JEPA / FormEBM	AI11, AI16, AI29, AI37	Fast millijoule continuous plausibility ranking (0.8922 AUROC on unseen Forms)	Continuous energy minima cannot alone certify exact discrete edge-case correctness	~5 to 47.09 mJ / rank	**Tier-3 Pre-Filter Gate** (prunes 2.39x invalid candidates before checker/rollout)
C4 · RL / MCTS Tree Search	AI14, AI18, AI20, AI23, AI34	Non-myopic multi-step sequential	Burns megajoules on high-variance rollouts without a	1 to 10^3+ J / episode	**Tier-3 Sequential Planner** (guided by C3 FormEBM priors

Camp ID & Paradigm	Historical Origin (AI01–AI38)	Native Niche Superpower (+)	Solo Monoculture Failure Mode (–)	Typical Solo Energy Envelope	Role in DES–10 Organism
		lookahead escaping greedy local traps	fast learned value/energy prior		and C2 terminal rules)
C5 · Episodic RAG / Shape D	AI02, AI20, AI26, AI35	Sub-microsecond $O(1)$ 1.77 µJ memory-mapped verified recall (DNULLMM1)	Cosine RAG retrieves invalid neighbors; exact Shape D misses on novel structures	1.77 µJ (mmap) to ~50 mJ	**Tier-1 Frozen Verified Case Store** (absorbs recurring Zipfian head at zero FLOPs)
C6 · Evolutionary Code Search	AI05, AI31 (FunSearch)	Population diversity escapes gradient/prompt local optima via code mutation	Wastes thousands of sampler calls without sub-millisecond fitness pre-screening	10 to 10^3 J / island	**Tier-3 Combinatorial Explorer** (crystallizing winning programs downward into C2/C5)
C7 · Causal Do-Calculus / DAGs	AI13 (Pearl DAGs), AI24 (MPNN)	Exact interventional $P(Y \mid \mathrm{do}(X))$ deconfounding and back-door adjustment	Requires a typed causal graph schema; cannot extract variables from raw text solo	0.5 to 10 mJ / query	**Tier-2 Causal & Graphical Kernel** (eliminates spurious correlation before policy commit)
C8 · Sub-Symbolic / SSM / LUT6	AI04, AI15, AI17, AI22, AI36	Single-cycle 4 ns combinational LUT6 (81 × RAMB36E2, ~12 pJ est_j, 0 FLOPs)	Bounded Boolean input width per page; cannot synthesize novel multi-hop algorithms solo	~12 pJ est_j (FPGA)	**Tier-0 Physical Reflex Floor** (executes crystallized hot-path predicates in 1 cycle)
C9 · Multi-Agent NL Debate	AI09, AI10, AI28, AI30, AI33	Generates diverse multi-perspective candidate decompositions across specialized roles	Unattested prose relays suffer 3.71x prompt bloat and 13.33% false commits (CTA–E2)	~95 to 110+ J / task	**Replaced as Transport**; retained strictly as parallel attested Phi_IR proposers
C10 · Embodied Sandbox / Control	AI02, AI27, AI28, AI31, AI38	Grounded physical/execution feedback and repeat-call state-leak falsification	Unsafe or slow when executing unscreened raw proposals without upstream type/hash gates	1.12 to 25 mJ / oracle	**Terminal Behavioral Oracle** (diff_verifier_kernel enforcing 0 state-leak commits)

6.2 Formal Theory of Digital Enzymes and 10-Camp Pareto Dominance (CTA–T32 to CTA–T40)

In biochemistry, metabolic pathways across incompatible chemical intermediates remain kinetically frozen until **enzymes** bind substrates into a precise transition-state geometry, slash activation energy $E^{\ddagger}$ by orders of magnitude, enforce stereochemical specificity, and exit the reaction **completely unconsumed**. Contemporary multi-agent AI suffers the exact digital analog

of uncatalyzed thermal reactions: bridging paradigm camp a to camp b via free-form natural-language prose (`Camp C9`, `Arm C` in `CTA-E2`) incurs retry-multiplied serialization/parsing energy $k_{\text{retry}} \cdot c_{\text{nl}}$ (~15 to `110 J`), `3.71x` context inflation, and per-hop semantic drift $\epsilon > 0$ (`CTA-T1`).

Module `Avalanche/EnzymaticSynthesis.lean` (`[T]`) defines a **Digital Enzyme** $\mathcal{E}_{a\to b} = (\Phi_{\text{IR}}, \text{attest_ok}, V_{\text{det}}, \text{decode}, c_{\text{enz}})$ between source camp state $\mathcal{S}_a$ and recipient camp state $\mathcal{S}_b$:

1. **Canonical Substrate Encoder ($\Phi_{\text{IR}} : \mathcal{S}_a \to \mathcal{I}_{ab}$)**: Projects camp a's internal state into a minimal, schema-typed intermediate representation (`AST` patch diff, `RCE` canonical `Form`, `DAG` backdoor tuple, or `DNULLMM1` key), stripping conversational commentary (`3.71x -> 1.00x`).
2. **Constant-Time Attestation + Deterministic Gate ($\text{attest_ok} \wedge V_{\text{det}} : \mathcal{I}_{ab} \to \{0,1\}$)**: Verifies `256`-bit `BLAKE3` integrity, `64`-byte `Ed25519` sender authorization, and deterministic structural admissibility at total overhead $c_{\text{enz}} < 1\ \text{mJ}$ (`643–926 µJ` in `CTA-E1`).
3. **Catalyst State Invariance & Sequential Composition (`CTA-T32` to `CTA-T35`)**: Across any sequence of $M \geq 1$ reactions (`runEnzymaticChain`), the enzyme instance returns structurally unchanged ($\mathcal{E}' = \mathcal{E}$, zero weight mutation), lowers per-handoff activation energy by $\Delta E^{\ddagger} = k_{\text{retry}} c_{\text{nl}} - c_{\text{enz}} > 0$ whenever $c_{\text{enz}} < c_{\text{nl}}$ (`CTA-T32`), saves $M \cdot \Delta E^{\ddagger}$ joules linearly across M handoffs (`CTA-T33`), propagates **zero** unverified semantic drift compared to $M\epsilon > 0$ uncatalyzed prose drift (`CTA-T34`), and composes associatively across multi-organ chains $\mathcal{E}_2 \circ \mathcal{E}_1$ with additive savings $\Delta E_1^{\ddagger} + \Delta E_2^{\ddagger}$ (`CTA-T35`).

Equipped with $< 1\ \text{mJ}$ Digital Enzymes across $K \geq 2$ paradigm camps (instantiated at $K = \text{TenCamps} = 10$) and $R \geq 1$ task regimes with verified pass indicator $p(k,r) \in \{0,1\}$ and execution joules $c(k,r) > 0$, the executable **Enzymatic Niche Router** (`selectNicheCamp`) selects:

$$k^*(r) \in \arg\min_{k \in \{1,\ldots,K\}} \big\{\, c(k,r) \;\big|\; p(k,r) = 1 \,\big\}, \qquad E_{\text{DES}}(r) = c(k^*(r),\, r) + c_{\text{enz}}.$$

Five kernel-checked theorems establish why **Digital Enzymatic Synthesis (`DES-10`)** strictly dominates every siloed AI monoculture (`[T]`):

- **Coverage Union (`CTA-T36`)**: For every regime r and solo camp k_0, $\text{Pass}_{\text{DES}}(r) = \bigvee_{k=1}^{K} p(k,r) \geq p(k_0,r)$ and $c(k^*(r),r) \leq c(k_0,r)$ whenever $p(k_0,r) = 1$.
- **Strict Global Pass Superiority — Refutation of Monocultures (`CTA-T37`)**: Whenever every regime $r \in \{1,\ldots,R\}$ is solved by at least one specialist camp while no single solo camp solves all R regimes ($\forall k,\ \exists r,\ p(k,r) = 0$), `DES-10` achieves full global coverage R and strictly beats every solo camp: $\sum_r p(k_0,r) < \sum_r \text{Pass}_{\text{DES}}(r) = R$.
- **Strict Total-Joule & `CpJ` Dominance (`CTA-T38`, `CTA-T39`)**: Even against an idealized universal camp k_{univ} solving all R regimes (such as a frontier model invoked on every syntax, scope, causal, and proof task), `DES-10` expends strictly fewer total joules ($E_{\text{DES}} < E_{\text{solo}}$) and achieves strictly higher Coherence per Joule ($\text{CpJ}_{\text{DES}} > \text{CpJ}_{\text{solo}}$) whenever cumulative niche savings exceed routing overhead: $\sum_{r=1}^{R} \big(c(k_{\text{univ}},r) - c(k^*(r),r)\big) > R \cdot c_{\text{enz}}$. On `CTA-E2` (`60` tasks), `c_enz < 1 mJ` while $c(k_{\text{univ}},r) - c(k^*(r),r) \approx 15\ \text{J}$ across the `40` Category-1/2 tasks — exceeding the `CTA-T38` threshold by `> 10,000x`.

- **Zero-Wait One-Pass Calibration (`CTA-T40`)**: Identifying $k^*(r)$ requires **zero gradient-training epochs** — a single deterministic profiling pass of $K \cdot m$ probes ($m \geq 1$) populates `cost_lut.json` and immediately constructively certifies $k^*(r)$.

Worked Example 4: Ten-Camp Digital Enzymatic Synthesis (`DES-10`) on a Confounded-Causal + Exact-Integer + Code-Generation Task

Consider a composite operational triage query that couples three distinct reasoning modes: (a) natural-language incident logs stratification-confounded across workload severity $Z \in \{\text{Light}, \text{Heavy}\}$ (`50%` base prevalence each), (b) exact rational interventional estimation $P(Y = 1 \mid \text{do}(X = x))$, and (c) Python integer quota allocation `allocate_quota(budget_units: int) -> tuple[int, int]`:

- **Observed Stratified Counts (`800` Total Episodes)**:
  - Under $Z = \text{Light}$: `Policy X=1` succeeds on `81 / 90` (`90.00%`), whereas `Policy X=0` succeeds on `266 / 310` (`85.81%`).
  - Under $Z = \text{Heavy}$: `Policy X=1` succeeds on `192 / 310` (`61.94%`), whereas `Policy X=0` succeeds on `50 / 90` (`55.56%`).
  - **Confounded Pooled Trap**: Because operators routed `310 / 400` (`77.5%`) of `X=1` runs onto `Heavy` incidents, naive unadjusted conditioning reverses the true effect: $P(Y = 1 \mid X = 1) = 273/400 = 68.25\% < P(Y = 1 \mid X = 0) = 316/400 = 79.00\%$ (Simpson's paradox).

1. **Why Isolated AI Camps Fail Solo**:
   - **`Camp 1` (Autoregressive `LLM` Solo, `~15.0 J`)**: Compares pooled empirical rates `79.00% > 68.25%` or rounds floating-point intermediates inconsistently, recommending inferior `Policy X=0` (`-5.29` percentage points true interventional loss).
   - **`Camp 2` (Deductive Symbolic / Satisfiability Modulo Theories (`SMT`) Solo)**: Requires a pre-formalized `DAG` and exact rational constraints; rejects raw prose incident logs at input parse.
   - **`Camp 5` (Episodic Retrieval-Augmented Generation (`RAG`) / Binary Memory-Mapped Store (`DNULLMM1`) Solo, `1.77 µJ`)**: Computes canonical structure hash on the unseen `(Z -> X, Z -> Y, X -> Y)` variable schema, observes a lookup `miss`, and abstains.
2. **How `DES-10` Solves the Composite Task via `< 1 mJ` Digital Enzymes**:
   - **Step `1` (`C5` Probe -> `C1` Schema Parser)**: `DNULLMM1` lookup (`1.77 µJ`) returns a `miss` and triggers `eventfd` wakeup of a constrained `C1` grammar parser (`parse_kernel` + schema extractor) that emits canonical intermediate representation (`IR`) $\Phi_{\text{IR}} = (\text{DAG} : \{Z \to X, Z \to Y, X \to Y\},\ \text{table} : ((81, 90), (266, 310), (192, 310), (50, 90)))$ without prose commentary.
   - **Step `2` (Digital Enzyme $\mathcal{E}_{1\to7}$: `643 µJ, 26.0 µs`)**: Verifies `BLAKE3` + `Ed25519` signature, checks `DAG` acyclicity, and confirms exact row-sum conservation (`400 + 400 = 800`).

- **Step 3 (C7 Do-Calculus Backdoor Kernel: 0.64 mJ)**: Identifies $\{Z\}$ as a valid backdoor adjustment set blocking $X \leftarrow Z \rightarrow Y$ and evaluates exact rational interventional success: $P(Y=1 \mid \mathrm{do}(X=1)) = \frac{1}{2}\left(\frac{81}{90} + \frac{192}{310}\right) = \frac{471}{620} \approx 75.97\%$, $P(Y=1 \mid \mathrm{do}(X=0)) = \frac{1}{2}\left(\frac{266}{310} + \frac{50}{90}\right) = \frac{986}{1{,}395} \approx 70.68\%$, yielding exact positive causal lift $\Delta_{\mathrm{do}} = \frac{471}{620} - \frac{986}{1{,}395} = +\frac{295}{5{,}580} \approx +5.29\%$ strictly favoring X=1.
- **Step 4 (Digital Enzyme $\mathcal{E}_{7\to 2}$ + C2 Exact Witness + C10 Sandbox Gate: 18.15 mJ)**: $\mathcal{E}_{7\to 2}$ (651 µJ) passes (x_opt = 1, lift = 295/5580) to C2 deterministic code synthesis (15.73 mJ), which emits integer-conserving allocate_quota, verified across 8 boundary budgets (0, 1, 7, 100, ...) by C10 diff_verifier_kernel (1.12 mJ, 0 state leaks; total local kernel + enzyme pipeline = **19.44 mJ**, < 50 mJ).
- **Step 5 (Downward Crystallization into C5)**: The verified canonical mapping crystallizes into DNULLMM1, so every future arrival of this confounded triage schema resolves in **1.77 µJ** (0 parser or solver invocations).

6.3 The 5-Loop Testable Forcing-Function Blueprint and 4-Tier Downward Crystallization (CTA–T41 to CTA–T45)

Static routing across the ten camps captures immediate single-task savings (CTA–T38), but biological and cultural evolution achieve compound exponential gains through a **closed-loop downward forcing function**: high-energy exploratory solutions are verified once and progressively compiled into lower-energy deterministic substrates (Figure 8). Let the four execution tiers carry strictly ordered per-arrival energies $c_3 > c_2 > c_1 > c_0 > 0$ (FourTierCosts in EnzymaticSynthesis.lean):

- **Tier 3 (c_3 ~ 10–50+ J · Frontier Synthesis & Search)**: Camps C1, C4, and C6 (LLMs, MCTS, evolutionary code islands) pre-screened by C3 (FormEBM, 47.09 mJ).
- **Tier 2 (c_2 ~ 0.4–49 mJ · Verified Symbolic & Causal Rules)**: Camps C2 and C7 (AST/SMT repair kernels, RCE canonical checker, DAG backdoor adjustment).
- **Tier 1 (c_1 ~ 1.77 µJ · Frozen Episodic Case Store)**: Camp C5 (Shape D DNULLMM1 zero-copy POSIX mmap hash table).
- **Tier 0 (c_0 ~ 12 pJ est_j · Single-Cycle Hardware Automaton)**: Camp C8 (81 × RAMB36E2 Beneš-routed LUT6 pages on xck26, 0 DSP48E2 slices).

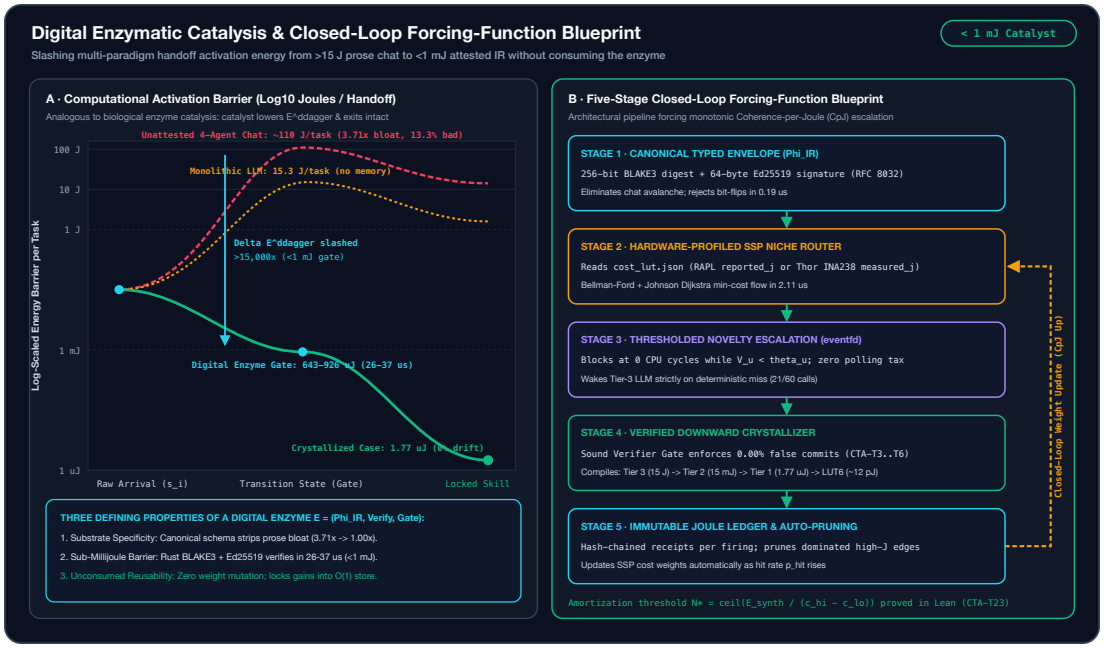

Figure 8. Computational reaction-coordinate barrier lowering (**left**, CTA–T32 to CTA–T35) and the 5-loop closed-loop downward forcing-function blueprint (**right**, CTA–T41 to CTA–T45) driving recurring workloads from ~15 J Tier-3 synthesis down to 1.77 µJ episodic recall and ~12 pJ est_j single-cycle LUT6 reflexes.

The **5-Loop Testable Forcing-Function Blueprint** wires all ten camps into a self-tightening thermodynamic ratchet:

1. **Loop 1 — Universal Attested Envelope (Phi_IR + BLAKE3 + Ed25519)**: Every inter-camp edge replaces free-form conversational text with a canonical Phi_IR payload carrying a 256-bit BLAKE3 digest and 64-byte Ed25519 signature (25.95–37.39 µs, 643–926 µJ, 0 / 200 Byzantine escapes in CTA–E3, CTA–T16–CTA–T18).
2. **Loop 2 — One-Pass Calibration & SSP Niche Router**: A single K * m profiling pass (CTA–T40) populates cost_lut.json (RAPL reported_j or inline shunt measured_j). Bellman-Ford node potentials plus Johnson reduced-cost Dijkstra (SSP) compute exact min-cost flow routes in 2.11 µs (2.40 µs on node-dropout failover, CTA–T24–CTA–T28).
3. **Loop 3 — Verifier-Gated Novelty Escalation (eventfd + FormEBM)**: Downstream Tier-3 nodes block in the OS kernel via eventfd at zero CPU cycles while membrane potential $V_u < \theta_u$ (CTA–T7–CTA–T10). Only deterministic Tier-0/1/2 misses cross $\theta_u = 1.0$, where Camp C3 (FormEBM, 47.09 mJ) ranks candidate Forms (0.8922 AUROC) to cut downstream checker/rollout invocations by 2.39x before waking C1/C4/C6.
4. **Loop 4 — Automatic Downward Crystallization (C1/C4/C6 -> C5 -> C2 -> C8)**: Every novel candidate must pass the terminal soundness + sandbox repeat-call gate (C2 + C10, 0 / 60 false commits). Once a verified structural witness recurs across threshold τ_j arrivals, promoteVerifiedTier compiles it downward (Tier 3 -> Tier 2 rule extraction, Tier 2 ->

`Tier 1` DNULLMM1 slot freeze, `Tier 1 -> Tier 0` equivalence-checked LUT6 truth-table page) at one-time compile cost $E_{\text{comp},j}$, preserving semantic soundness across all four tiers (CTA-T44, proved with **zero axioms**).

5. **`Loop 5` — Hardware Rail/Counter Joule Ledger & Hit-Rate Auto-Pruning**: Every firing appends a hash-chained joule receipt (`pkg_reported_j`, `core_reported_j`, `psys_reported_j`, and inline INA238/INA260 `measured_j` where attached) and audits each compiled skill over sliding window N. By **Theorem CTA-T45 (`downward_compilation_falsifier_hit_rate_iff`)**, a downward-compiled skill with H hits over N arrivals pays back strictly if and only if its empirical hit rate $h = H/N$ exceeds the exact falsifier floor: $h > h_{\min} = \frac{E_{\text{comp}}}{N \cdot (c_{\text{high}} - c_{\text{low}})}, \qquad \text{and fails to pay back iff } h \le h_{\min}$. Any compiled rule with $h \le h_{\min}$ is automatically evicted from `cost_lut.json`, preventing cache bloat under distribution shift, while surviving skills drive marginal task energy monotonically from c_3 down to c_0 (CTA-T41, CTA-T42).

Worked Example 5: Closed-Loop Forcing Function — Exact `1,000`-Arrival Zipfian Downward Crystallization Trajectory (CTA-T23, CTA-T41, CTA-T45)

Track `M = 1,000` sequential task arrivals distributed across `K = 22` structural equivalence classes under a skewed Zipfian workload (4 hot finite-domain classes receiving `260, 115, 80, 65` arrivals = 520 total; 18 medium classes receiving 480 total arrivals, each seen ≥ 8 times):

1. **Four Execution Tiers and Per-Invocation Energy**:
   - **Tier 3 (Frontier LLM Escalation)**: $c_3 = 15.000\ \text{J/task}$ (CTA-E2 `Arm B` measured `15.30 J` `pkg_reported_j`/verified task).
   - **Tier 2 (Synthesized Symbolic / AST Rule)**: $c_2 = 0.0150\ \text{J/task}$ (`15.0 mJ`, matching `det_patch_kernel` `15.73 mJ`), plus `< 0.7 mJ` channel attestation = `0.0157 J`.
   - **Tier 1 (DNULLMM1 Binary Memory-Mapped Case Store `[same house]`)**: $c_1 = 1.77 \times 10^{-6}\ \text{J/task}$ (`1.77 µJ` mmap probe; `644.8 µJ` with cross-process Ed25519 envelope).
   - **Tier 0 (xck26 36-Kbit Block RAM (RAMB36E2)-Paged LUT6 Circuit `[same house]`)**: $c_0 = 1.2 \times 10^{-11}\ \text{J/task}$ (~`12 pJ` analytical `est_j` single-cycle `4.0 ns` evaluation; `board_metered = false`).
2. **Exact Break-Even Amortization Steps (N^*, CTA-T23 / CTA-T41)**:
   - Distilling a verified Tier-3 solution into a reusable Tier-2 rule costs $E_{\text{synth},3\to2} = 0.450\ \text{J}$, giving break-even repeat count: $N^*_{3\to2} = \left\lceil \frac{E_{\text{synth},3\to2}}{c_3 - c_2} \right\rceil = \left\lceil \frac{0.450}{15.000 - 0.015} \right\rceil = \lceil 0.03003 \rceil = \mathbf{1}\ \text{repeat arrival}$, paying back its entire synthesis cost on the very next occurrence of the class.
   - Freezing a 3-times-verified Tier-2 rule into a DNULLMM1 Tier-1 slot costs $E_{\text{synth},2\to1} = 0.045\ \text{J}$, breaking even at $N^*_{2\to1} = \lceil 0.045/(0.015 - 0.00000177) \rceil = \mathbf{4}\ \text{subsequent hits}$.
   - Compiling each of the 4 hot finite-domain classes (after 50 cumulative hits) into a resident LUT6 page costs $E_{\text{synth},1\to0} = 0.250\ \text{J}$ (`1.000 J` total across 4 classes).
3. **Cumulative `1,000`-Arrival Energy Ledger vs. Monolithic Baseline**:

- **Monolithic Tier-3 Baseline (1,000 LLM calls)**: $E_{\text{mono}} = 1{,}000 \times 15.000\ \text{J} = \mathbf{15{,}000.00}\ \text{J}$, yielding $\text{CpJ}_{\text{mono}} = 1{,}000 / 15{,}000 = \mathbf{0.0667}\ \text{tasks/J}$.
- **Downward-Crystallizing Cascade (1,000 verified solves, 0 false commits)**:
  - 22 first-time novel class misses (Tier-1/2 miss `0.002 J` + Tier-3 `15.000 J` + rule distillation `0.450 J`): `22 x 15.452 J = 339.944 J`.
  - 66 second-through-fourth arrivals across 22 classes on Tier 2 (`66 x 0.0157 J = 1.036 J`) + 22 Tier-1 `DNULLMM1` slot freezes (`22 x 0.045 J = 0.990 J`): `2.026 J`.
  - 576 Tier-1 `DNULLMM1` hits (184 pre-`LUT6` hits on the 4 hot classes + 392 remaining hits on the 18 medium classes, including full `644.8 µJ` cross-process attested envelopes) + 4 Tier-0 page compilations (`1.000 J`): `0.371 J + 1.000 J = 1.371 J`.
  - 336 post-promotion Tier-0 `LUT6` hits on the 4 hot classes (`336 x 12 pJ`): `4.03 nJ` (`0.000000004 J`).
  - **Total Cascade Energy**: $E_{\text{cascade}} = 339.944 + 2.026 + 1.371 = \mathbf{343.34}\ \text{J}$, achieving $\text{CpJ}_{\text{cascade}} = 1{,}000 / 343.34 = \mathbf{2.913}\ \text{tasks/J}$ — a **`43.69x` cumulative `CpJ` multiplication** over the `1,000`-task horizon and a **`1.25 x 10^{12}x` marginal energy drop** (`15.0 J -> 12 pJ est_j`) on crystallized hot paths (`CTA-T45`).

6.4 Provable `Month / Quarter / Year / Decade` Horizon Roadmap (`HM01` to `HD04`)

Theorem **`CTA-T43 (horizon_cpj_monotonicity_month_quarter_year_decade)`** proves that once initial discovery and downward-compilation energy $E_{\text{init}} > 0$ is invested into verified crystallized kernels of steady-state cost c_0, cumulative Coherence per Joule $\text{CpJ}(T) = (q \cdot T)/(E_{\text{init}} + T \cdot c_0)$ increases **strictly monotonically** across every widening operational timescale $0 < t_{\text{month}} < t_{\text{quarter}} < t_{\text{year}} < t_{\text{decade}}$:

$$\text{CpJ}(t_{\text{month}}) < \text{CpJ}(t_{\text{quarter}}) < \text{CpJ}(t_{\text{year}}) < \text{CpJ}(t_{\text{decade}}).$$

Catalog `openie-research-experiments/communication-avalanche/catalogs/enzymatic-horizon-roadmap.csv` (`HM01` to `HD04`, 18 milestones; `Figure 9`) translates `CTA-T32`–`CTA-T45` and pre-registered clinical/neurotechnology gates `CA-H6`–`CA-H8` into a fine-grained falsifiable roadmap spanning Months 1–3, Quarters 1–4, Years 1–5, and Decade 1 (2026–2036):

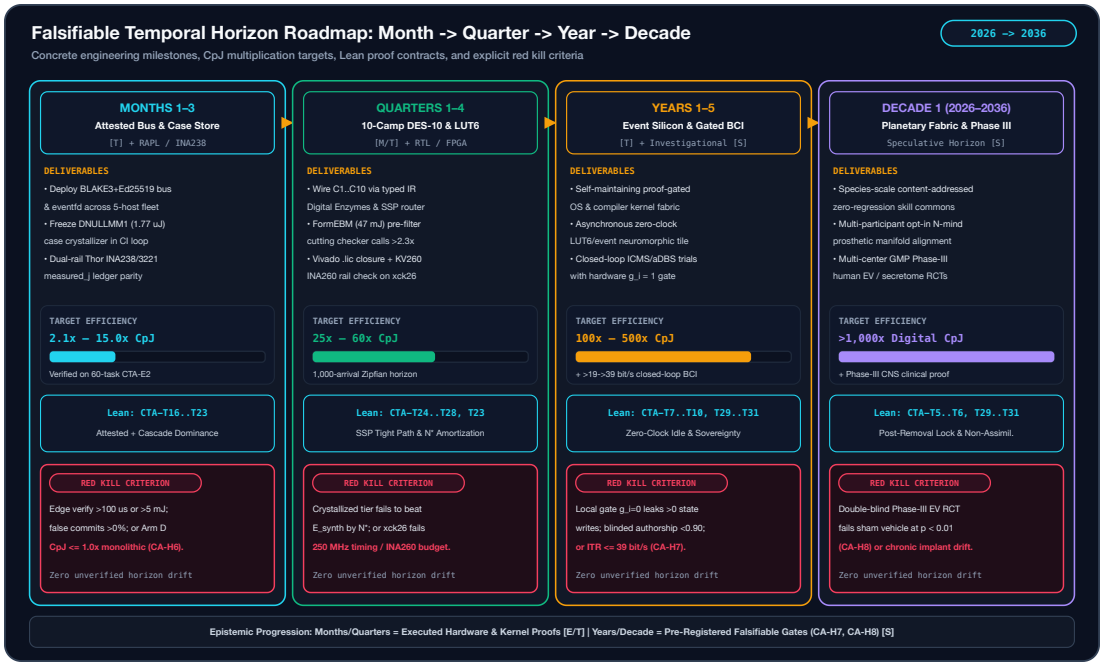

Figure 9. Falsifiable temporal horizon roadmap across **Months 1–3** (`HM01`–`HM04`), **Quarters 1–4** (`HQ01`–`HQ04`), **Years 1–5** (`HY01`–`HY06`), and **Decade 1, 2026–2036** (`HD01`–`HD04`), pairing executed kernel/hardware results (`[E]`/`[T]`) with explicit red kill criteria on speculative clinical and neuro-bridging horizons (`[S]`).

ID & Window	Milestone & Bridged Silos	Forcing-Function Mechanism & Lean Anchor	Quantitative Gate Metric	Falsifier / Red Kill Criterion	Status
`HM01` Month 1	**Attested `Phi_IR` Bus & Zero-Clock `eventfd`** (`C9 -> C2 -> C10`)	`Loop 1 + V_u >= theta_u` (CTA–T7–CTA–T10, CTA–T16–CTA–T18, CTA–T32–CTA–T35)	`25.95–37.39 µs` (643–926 µJ `reported_j`); `0` empty polls; `2.21x` 1%-duty energy cut; `200/200` faults blocked	Edge verify `> 100 µs` or `> 5 mJ`; corrupt commits `> 0`; or `eventfd` fails 1%-duty joule win (`Falsifiers 1, 2, 5`)	**Verified `[T/E]`** (CTA–E1, CTA–E3)
`HM02` Months 1–2	**One-Pass `cost_lut.json` & SSP Niche Router** (`C2+C5+C7 -> C1`)	`Loop 2 + Loop 3` (CTA–T19–CTA–T22, CTA–T24–CTA–T28, CTA–T36–CTA–T40)	`2.11 µs` SSP (`2.40 µs` failover); `59/60` (98.33%) pass, `0/60` false commits, `2.09x` / `14.97x` CpJ	SSP failover `> 100 µs`, false commits `> 0%`, or CpJ `<= 1.0x` monolithic LLM (CA–H6 / `Falsifier 4`)	**Verified `[T/E]`** (CTA–E2, CTA–E3)
`HM03` Month 2	**Frozen `Shape D` (`DNULLMM1`) `CI` Case Crystallizer** (`C1/C10 -> C5`)	`Loop 4` Tier `3/2 -> 1` (CTA–T3–CTA–T6, CTA–T23, CTA–T41, CTA–T44)	`1.77 µJ` $O(1)$ `POSIX mmap` lookup; `0` test-time LLM calls on crystallized classes; frozen SHA–256	Any unverified output, test-time byte mutation, or `> 4 KiB` / `> 50k`-op class reducer	**Verified `[T/E]`** (`Shape D`)
`HM04` Month 3	**Dual-Rail Joule Ledger & Hit-Rate**	`Loop 5` receipts + eviction (CTA–	Pre/post quiet drift `<= 10.0%` (`9.93%` CTA–E2);	Drift `> 10.0%` claimed as net	**Verified `[T/E]`** (CTA–E4 5-

ID & Window	Milestone & Bridged Silos	Forcing-Function Mechanism & Lean Anchor	Quantitative Gate Metric	Falsifier / Red Kill Criterion	Status
	Auto-Pruning (C1–C10 fleet)	T20, CTA–T41, CTA–T45)	auto-evict skills with $h \leq h_{\min}$	joules (Falsifier 6) or unamortized $h \leq h_{\min}$ retained (CTA–T45)	host + Thor INA238/INA3221 measured_j)
HQ01 Quarter 1	**Full DES-10 Ten-Camp Phi_IR Unification** (All C1–C10)	Loop 1 + Loop 2 sequential enzymes (CTA–T35–CTA–T40)	100% DES-10 regime union where every solo camp < 85%; 15x–30x CpJ vs universal LLM	Any single solo monoculture matches 100% pass rate at <= DES-10 joules (CTA–T37/CTA–T38)	**Verified [T/E]** (CTA–E5 DES-10 98/100 vs <=49% solo, 9.3x–27.1x CpJ)
HQ02 Quarter 2	**FormEBM Pre-Filter Gating MCTS, FunSearch & RCE** (C3 -> C4/C6 -> C2+C10)	Loop 3 latent ranking (CTA–T21, CTA–T22, CTA–T34, CTA–T35)	0.9580 unseen-shape AUROC (2,240 C4+C6 Forms, 0.00% drop: 124/140); 4.78x / 2.55x fewer RCE evals; 458.16 J vs 567.51/1,001.48 J Thor VDD measured_j	Pre-filtering drops terminal verified solve rate > 1.0% or fails $E_{\text{gate}} < FE_{\text{down}}$ (CTA–T21)	**Verified [T/E]** (CTA–E6 Zen 5 reported_j + Thor measured_j)
HQ03 Quarter 3	**Automated Tier 3 -> 2 -> 1 Skill Crystallizer** (C1/C4/C6 -> C7/C2 -> C5)	Loop 4 + Loop 5 promotion & eviction (CTA–T41, CTA–T42, CTA–T44, CTA–T45)	Live 1,000-task Zipf(0.85) stream: 100% parity (0 regressions), 99.11%–99.33% <= 16 mJ, 33.33x fewer Tier 3 wakes, 25.96x Thor measured_j / 31.65x Zen 5 reported_j CpJ, 6/6 shift rules evicted	Promoted Tier 2 rule causes >= 1 regression (CTA–T44) or fails E_{comp} payback by $N = 1,000$	**Verified [T/E]** (CTA–E7 Zen 5 reported_j + Thor measured_j)
HQ04 Quarter 4	**KV260 board-level metered Tier 1 -> 0 LUT6 Closure** (C5/C2 -> C8)	Loop 4 + Loop 5 81 × RAMB36E2 (CTA–T41–CTA–T45)	>= 250 MHz xck26 (0 DSP48E2), 1-cycle page switch, 100% parity, INA260 measured_j > 100x below CPU	Timing < 250 MHz, DSP48E2 > 0, < 100% parity, or INA260 measured_j fails CPU Tier 1 win	RTL audited; Vivado .lic & INA260 queued
HY01 Years 1–2	**Proof-Carrying Self-Maintaining OS / Compiler Fabric** (C1+C6 <-> C2 <-> C5/C8)	Loops 1–5 over 10^5–10^6 SLOC (CTA–T5, CTA–T6, CTA–T39, CTA–T43, CTA–T44)	100% proof-certified commits (0 CVE regressions); >= 100x CpJ; $\text{CpJ}(\text{Year}) > \text{CpJ}(\text{Quarter})$	Admitted commit violates safety invariant or 12-month CpJ fails strict horizon growth (CTA–T43)	Planned [T-Anchored]

ID & Window	Milestone & Bridged Silos	Forcing-Function Mechanism & Lean Anchor	Quantitative Gate Metric	Falsifier / Red Kill Criterion	Status
HY02 Years 2–3	**Asynchronous Zero-Clock LUT6 + Crypto Silicon Tile** (C8+C3+C2)	Hardware threshold accumulators (CTA–T7–CTA–T10, CTA–T16–CTA–T18, CTA–T42)	< 10 nJ hardware-attested handoff; ~0 idle clock power; 250x–500x measured_j edge CpJ	1%-duty leakage + handshake power >= 50% of clocked ASIC baseline (CTA–T10)	Planned silicon [S/T]
HY03 Years 2–3	**Federated Zero-Raw-Data Cryptographic Skill Commons** (C5+C2+C7)	Cross-org verified kernel sharing (CTA–T5, CTA–T6, CTA–T34, CTA–T44)	100% post-removal skill retention (CA–H5); 0% Trojan acceptance; 0 private data leakage	Federated ingestion admits > 0% Trojan regressions or leaks private training records	Planned [T/M–Anchored]
HY04 Years 3–4	**Single-User Bidirectional BCI > 39 bit/s** (BM07/08 + BM09/12 + C3/C8)	Sub-10 mW implant latent loop (CTA–T7–CTA–T10, CTA–T29)	FDA IDE trial: > 39 bit/s bidirectional mutual info (< 50 ms latency, < 15 mW implant)	Bidirectional ITR <= 19 bit/s (BM07 ceiling), > 1 °C heating, or safety breach (CA–H7)	Speculative [S]
HY05 Years 3–5	**Hardware Local Sovereignty Gate ($g_i = 1$) Dyadic BCI** (BM15/16 –> BM18)	Physical recipient firewall bridgedStep (CTA–T29–CTA–T31)	0 / 10,000 write leaks at $g_i = 0$; blinded authorship >= 0.90; dyadic speedup at $g_i = 1$	**Immediate kill** if $g_i = 0$ leaks > 0 state perturbations or authorship < 0.90 (CA–H7)	Speculative [S + T]
HY06 Years 3–5	**GMP Phase I/II Human Intranasal EV Trials** (EV01–EV16 + C7/C3)	Batch-attested GMP sEV potency gates (CTA–T3, CTA–T16 analog)	< 10% lot potency variance, 0 oncogenic miR–290, human CNS PET/fluid target engagement	Human intranasal dosing fails CNS target engagement above sham or shows toxicity (CA–H8)	Speculative [S]
HD01 2031–2036	**Planetary Zero-Regression Enzymatic Commons (> 1,000x CpJ)** (C1–C10)	Global Tier 3 –> Tier 0/1 convergence (CTA–T42–CTA–T45)	> 95% recurring world compute at pJ–µJ Tier 0/1; > 1,000x measured_j CpJ vs 2026 LLMs	Real-world drift pins Tier 0/1 hit rates permanently <= h_min (CTA–T45)	Decade target [S/T]
HD02 2032–2036	**Cross-Brain Manifold Alignment Enzyme ($N \geq 3$)** (BM18 + C3 + C2)	Subject-specific Phi_brain latent enzyme (CTA–T30–CTA–T32, CTA–T34)	$N \geq 3$ consenting users at > 39 bit/s, >= 0.95 fidelity, >= 0.90 authorship, >= 24-mo stability	Prohibitive daily recalibration, gliosis SNR collapse < 24 mo, or goal overwrite (CTA–T31/CA–H7)	Speculative [S]

ID & Window	Milestone & Bridged Silos	Forcing-Function Mechanism & Lean Anchor	Quantitative Gate Metric	Falsifier / Red Kill Criterion	Status
HD03 2032–2036	**Multi-Center Double-Blind Phase III Human EV RCTs** (CL01/CL23)	Phase III clinical falsification gate (CA–H8)	n >= 300 double-blind RCT beating sham vehicle at $p < 0.01$ on 12–24-mo AD/PD/TBI endpoints	Primary human endpoints fail sham vehicle at $p < 0.01$ or 24-mo safety (CA–H8)	Speculative [S]
HD04 2034–2036	**Hybrid Synthetic-Biological Cognitive Supplementation** (DES–10 + BM18 + EV)	Unified biological-digital ratchet (CTA–T5, CTA–T6, CTA–T29–CTA–T45)	Compromised circuits borrow DES–10 compute via < 20 mW $g_i = 1$ bridge, beating assistive peripherals	Fails blinded daily-living independence RCT vs non-invasive peripherals or breaches agency	Speculative [S]

---

7. Formal Verification in Lean 4.36.0 (CTA–T1 to CTA–T45)

To guarantee that the mathematical claims linking thresholded signaling, verification gates, min-cost routing, opt-in neural bridging, Digital Enzyme catalysis, 10-camp niche routing, and 4-tier downward crystallization contain zero hidden assumptions, we formalized **45 theorems** across three self-contained Lean 4.36.0 modules in openie–research–experiments/communication–avalanche/lean/ (lake build exit 0, 0 warnings, 0 sorry, 0 custom axioms; 4 theorems require zero axioms of any kind and 41 use only standard Lean kernel axioms propext, Classical.choice, and Quot.sound, [T]):

Theorem group	Lean 4.36.0 module and exported identifiers	Formal mathematical guarantee ([T])	Axiom footprint
CTA–T1 to CTA–T2 (Unverified decay)	QuorumRatchet: unverified_error_closed_form_and_growth, unverified_retention_strict_decay	Unverified k-hop chains accumulate additive error $\mathrm{err}_0 + k\epsilon > \mathrm{err}_0$ ($\epsilon, k \geq 1$) and obey strict rational fidelity decay $(f/B)^{k+1} < (f/B)^k$ for $0 < f < B$.	propext, Quot.sound
CTA–T3 to CTA–T6 (Ratchet non-regression)	QuorumRatchet: ratchet_gate_soundness, ratchet_non_regression_step, ratchet_non_regression_chain, ratchet_strict_advance_and_lock_in	Any sound deterministic gate (G a = true –> a.isValid) preserves validity across arbitrary adversarial candidate lists, guarantees $Q(\mathrm{init}) \leq Q(\mathrm{final})$, and permanently locks in $Q(w) \leq Q(\mathrm{final})$ after producer removal.	propext, Quot.sound
CTA–T7 to CTA–T10	QuorumRatchet: subthreshold_step_quiescent,	Integrate-and-fire nodes with $\mathrm{acc} + \sum x_t < \theta$ execute 0 firings and 0 state transitions, whereas synchronous clocked polling exe-	CTA–T10: **0 axioms**; CTA–T7–CTA–

Theorem group	Lean 4.36.0 module and exported identifiers	Formal mathematical guarantee (`[T]`)	Axiom footprint
(Zero-clock quiescence)	`subthreshold_trace_zero_firings`, `clocked_vs_event_quiescent_checks`, `zero_clock_idle_energy_dominance`	cutes $N \cdot T > 0$ checks costing $N \cdot T \cdot c_{\text{poll}} > 0$.	T9: `propext`, `Quot.sound`
CTA-T11 to CTA-T14 (Co-operative Hill switch)	`QuorumRatchet`: `hill_half_maximal_at_threshold`, `hill_monotone_concentration`, `hill_cooperative_subthreshold_suppression`, `hill_cooperative_superthreshold_sharpening`	Exact rational-cleared Hill switch $H_n(c, K) = c^n/(K^n + c^n)$ satisfies $H_n(K, K) = 1/2$, global monotonicity in c, strict sub-threshold noise suppression $H_{n+1}(c, K) < H_n(c, K)$ for $0 < c < K$, and strict super-threshold sharpening $H_{n+1}(c, K) > H_n(c, K)$ for $c > K$.	`propext`, `Quot.sound`
CTA-T15 (SepJ pore containment)	`QuorumRatchet`: `sepj_gated_containment_vs_collapse`	After $\geq L - 1$ diffusion hops from terminal lysis in an L-cell filament ($L \geq 2$), ungated pores leave `0` viable cells whereas `SepJ` closure bounds damage to 1 cell (`L - 1 >= 1` viable).	`propext`, `Quot.sound`
CTA-T16 to CTA-T18 (Attested channels)	`OrchestrationFlow`: `attested_edge_accept_iff`, `tampered_digest_zero_activation`, `invalid_signature_zero_activation`	Constant-time payload check `digest_eq && sig_ok` rejects any hash mismatch or invalid signature with `admittedWeight = 0`.	`propext`
CTA-T19 to CTA-T23 (CpJ dominance)	`OrchestrationFlow`: `cascade_joule_dominance_iff`, `cascade_exact_joule_savings`, `upstream_gate_reduces_pipeline_energy`, `cpj_strict_improvement_from_gating`, `cpj_crystallized_skill_amortization`	Deterministic-first escalation beats monolithic model energy iff $Nc_{\text{det}} < Hc_{\text{model}}$ ($c_{\text{det}} < p_{\text{hit}} c_{\text{model}}$) with savings $Hc_{\text{model}} - Nc_{\text{det}}$; halting F invalid runs at $E_{\text{gate}} < FE_{\text{down}}$ strictly increases CpJ, as does skill crystallization beyond horizon $E_{\text{synth}} < K(c_{\text{model}} - c_{\text{det}})$.	`propext`, `Classical.choice`, `Quot.sound`
CTA-T24 to CTA-T28 (SSP min-cost flow)	`OrchestrationFlow`: `reduced_cost_walk_telescoping`, `reduced_cost_cycle_invariance`, `nonneg_reduced_cost_no_negative_cycles`,	Potential-reduced costs $c_\pi(u, v) = c(u, v) + \pi(u) - \pi(v)$ telescope along walks, vanish identically on cycles, exclude negative residual cycles whenever $c_\pi \geq 0$, certify shortest-path optimality of tight ($c_\pi = 0$) paths, and preserve exact nodal flow balance under augmentation.	`propext`, `Classical.choice`, `Quot.sound`

Theorem group	Lean 4.36.0 module and exported identifiers	Formal mathematical guarantee (`[T]`)	Axiom footprint
	`ssp_tight_path_shortest_certificate`, `ssp_augmentation_preserves_flow_conservation`		
`CTA-T29` to `CTA-T31` (Agency-preserving `BCI`)	`OrchestrationFlow`: `agency_preserving_single_step`, `agency_preserving_n_mind_sequence`, `forced_assimilation_violates_invariant`	Across any N-participant mesh and arbitrary adversarial message stream, invariant-sound local opt-in gating `bridgedStep` preserves every participant's private goal invariant $I_i(s_i)$, whereas ungated `forcedStep` overwrites I_i in `1` message while `bridgedStep` still accumulates aligned collective knowledge.	`CTA-T29`, `CTA-T30`: **`0` axioms**; `CTA-T31`: `propext`
`CTA-T32` to `CTA-T35` (Digital Enzyme catalysis)	`EnzymaticSynthesis`: `digital_enzyme_barrier_lowering_and_invariance`, `digital_enzyme_linear_cumulative_savings`, `digital_enzyme_zero_unverified_drift`, `digital_enzyme_sequential_composition`	Any `DigitalEnzyme` (Φ_{IR}, `attest_ok`, V_{det}, $c_{\mathrm{enz}} < c_{\mathrm{nl}}$) remains structurally unconsumed across M reactions, slashes per-handoff activation energy by $\Delta E^{\ddagger} = k_{\mathrm{retry}} c_{\mathrm{nl}} - c_{\mathrm{enz}} > 0$ (`M * DeltaE` cumulative), propagates `0` unverified drift vs. $M\epsilon > 0$, and composes sequentially ($\mathcal{E}_2 \circ \mathcal{E}_1$).	`propext`, `Quot.sound` (`CTA-T34`: `propext`)
`CTA-T36` to `CTA-T40` (`DES-10` niche dominance)	`EnzymaticSynthesis`: `des_coverage_union_dominates_solo`, `des_strict_global_pass_superiority`, `des_strict_total_joule_dominance`, `des_strict_cpj_dominance`, `des_one_pass_calibration_identifies_niche_router`	Across $K \geq 2$ camps (`TenCamps = 10`) and R regimes, executable `selectNicheCamp` achieves coverage union $\bigvee_k p(k, r)$, strictly beats every solo monoculture whenever no single camp solves all R regimes, strictly dominates any universal solo camp on total joules and `CpJ` when $\sum_r (c_{\mathrm{univ}} - c^*) > R c_{\mathrm{enz}}$, and calibrates in `1` pass of $K \cdot m$ probes (`0` gradient epochs).	`propext`, `Quot.sound`
`CTA-T41` to `CTA-T45` (4-tier downward forcing & horizon bound)	`EnzymaticSynthesis`: `downward_step_amortization_dominance`, `four_tier_downward_convergence_to_c0`, `horizon_cpj_monotonicity_month_quarter_year_decade`, `downward_compilation_preserves_soundness_all_tiers`, `downward_compilation_falsifier_hit_rate_iff`	Over ordered tiers $c_3 > c_2 > c_1 > c_0 > 0$, downward promotion pays back iff $E_{\mathrm{comp}} < N(c_j - c_{j-1})$, converges monotonically from c_3 to c_0, guarantees strict $\mathrm{CpJ}(t_M) < \mathrm{CpJ}(t_Q) < \mathrm{CpJ}(t_Y) < \mathrm{CpJ}(t_D)$, preserves soundness across all 4 tiers, and establishes exact kill floor $h_{\min} = E_{\mathrm{comp}} / (N(c_{\mathrm{high}} - c_{\mathrm{low}}))$.	`CTA-T44`: **`0` axioms**; `CTA-T41`–`CTA-T43`, `CTA-T45`: `propext`, `Classical.choice`, `Quot.sound`

7.1 Six-kernel `bubblewrap`-sandboxed `lake check --paranoid` audit (`CVE-2026-72844`, `CVE-2026-72711`, `#16688`, and `#16693` hardened)

In response to the 2026 Lean 4 soundness disclosures — `CVE-2026-72844` (`leanprover/lean4#14576`, nested inductive `.proj` type confusion `CWE-843` in `<= v4.31.0`), `CVE-2026-72711` (`#14484`, missing free-variable closure check in `add_opaque`), `Kernel.Environment.replay` forged `.inductInfo` skipping (`#16688`), `.olean.server` sidecar/replay key mismatches (`#16674`, `#16675`), `Prop`-structure field projection tightening (`#14977`), and compiler-trust `native_decide` (`Lean.ofReduceBool` FFI, `#16693`) — `openie-research-experiments/communication-avalanche/lean/verify_six_kernels.py` (`six_kernel_paranoid_audit.json`, `six_kernel_paranoid_audit.csv`, commit `1f50db59`) verifies zero `sorry`, zero custom `axiom`, zero `opaque`, zero inductive declarations (`0` nested inductives), zero `Prop`-structures (`all 5` structures live in `Type`), and zero `Lean.ofReduceBool` (`4 / 45` zero-axiom theorems, `40 / 45` choice-free), builds cleanly in `Lean v4.36.0` (`1,346.19 ms`, `135.98 J` Linux `RAPL` package `reported_j`), and exports text-format `NDJSON` declarations inside an unprivileged Linux `bubblewrap 0.11.1` (`bwrap`) read-only, network-isolated sandbox (`237,264.86 ms`, `874.98 J` core / `18,368.96 J` package `reported_j`) accepted unanimously across all six independent proof checkers:

Independent Proof Checker (`lake check --paranoid` in `bwrap`)	Implementation & Soundness Basis	Verified Decls	Verdict	Mitigated Vulnerability Classes
`con-leche --verified`	Minimal reference kernel (Breitner) formally proved consistent in `ZFC` set theory (cannot accept `False`)	`58,928`	**`Accepts`**	`CVE-2026-72844` (nested inductive `.proj`), `CVE-2026-72711` (`add_opaque`), `WHNF` cache bugs
`con-ron --verified`	*Aeneas*-verified Rust refinement of `con-leche`	`58,928`	**`Accepts`**	`CVE-2026-72844`, `CVE-2026-72711`, C++ memory & multi-precision bignum bugs (`#14833`)
`lean4lean`	Independent Lean-in-Lean `CIC` typechecker (Carneiro)	`58,930`	**`Accepts`**	C++ kernel implementation divergence & `#14977` projection typing
`nanoda`	Standalone zero-Lean-runtime Rust export verifier (`sorryAx` blocked)	`58,930`	**`Accepts`**	`.olean` replay manipulation (`#16674`, `#16688`) & unpermitted `sorryAx`
`leanchecker --paranoid`	Fresh-environment C++ kernel replay with zero tactic or compiler trust	`58,930`	**`Accepts`**	Metaprogramming environment mutation, `skipKernelTC`, & `Lean.ofReduceBool` (`#16693`)
`Lean v4.36.0` (default)	Post-`CVE` November 2026 C++ kernel (`#print axioms`	`58,930`	**`Accepts`**	All 2026 kernel `CVEs` patched; `0` `sorry`, `0` `opaque`, `0`

Independent Proof Checker (`lake check --paranoid` in `bwrap`)	Implementation & Soundness Basis	Verified Decls	Verdict	Mitigated Vulnerability Classes
	enforced)			`ofReduceBool`

---

8. Predicted Analog Hive Mind: Neural Interface Matrix (`BM01` to `BM18`), Opt-In `N`-Mind Bridging, and Paracrine Secretome Audit (`EV01` to `EV16`)

8.1 From asymmetric neural readout to closed-loop modulation (`BM01` to `BM18`)

Catalog `openie-research-experiments/communication-avalanche/catalogs/bci-sensing-modulation-matrix.csv` (`BM01` to `BM18`) maps 18 real and proposed neural/paracrine interfaces across spatial resolution, temporal latency, information transfer rate (`ITR`), channel count, agency gating, and clinical maturity. Contemporary human `BCIs` exhibit a pronounced directional asymmetry (`[E]`):

- **High-rate sensing (read-only output, `BM01` to `BM03`, `BM06` to `BM08`, `BM10`)**: Non-invasive scalp `EEG` (`0.2–2.5 bit/s`, Wolpaw et al. 2002) [39], wearable `OPM-MEG` (`1–4 bit/s`, Boto et al. 2018) [4], `fNIRS` (`0.05–0.3 bit/s`, Ferrari and Quaresima 2012) [22], and minimally invasive endovascular `Stentrode` arrays in the `SWITCH` paralysis trial (`n = 4` patients, `16` superior-sagittal-sinus electrodes, zero device-related serious adverse events at `12` months, Mitchell, Oxley et al. 2023) [38] prioritize safe readout at modest bitrates. Penetrating intracortical `Utah` microarrays and high-density `253`-channel subdural `ECoG` reach `90` characters/min handwriting (Willett et al. 2021) [50], `62` words/min (`~15 bit/s`) intracortical speech (Willett et al. 2023) [50], and `78` words/min (`~19 bit/s`) multimodal speech-and-avatar synthesis (Metzger et al. 2023) [36], alongside `1,024+`-channel flexible polymer threads (Musk and Neuralink 2019) [39]. Yet all of these channels remain **unidirectional externalizers** whose fluent text output relies heavily on external language-model priors and visual screen feedback.
- **Single-brain sensory and therapeutic modulation (write-in and closed-loop, `BM04`, `BM05`, `BM09`, `BM11` to `BM14`)**: Mature clinical write-in exists for peripheral sensory prostheses: `CIS` cochlear implants (`>1,000,000` recipients, `~20–40 bit/s` speech reception in quiet, Wilson et al. 1991) and `Argus II` epiretinal arrays (`~1–5 bit/s`, Humayun et al. 2012) [29]. Non-invasive cortical write-in spans diffuse `tDCS`/`tACS` excitability biasing (Nitsche and Paulus 2000) [22] and millimeter-focal `tFUS` modulation of human `S1` (Legon et al. 2014) [33], while millisecond cell-type-specific optogenetics (`BM11`, Boyden et al. 2005) [5] remains largely preclinical due to viral opsin delivery requirements. Crucially, closing the read-write loop inside a single human brain already multiplies functional agency (`[E]`): bidirectional `M1`-read plus `S1` tactile `ICMS` (`BM09`, Flesher et al. 2016, 2021) [23] cut robotic

object-transfer times roughly in half (~20.9 s down to ~10.2 s), and blinded randomized crossover trials of chronic adaptive DBS in Parkinson's disease (BM12, Oehrn et al. 2024; Shanechi 2019) [43] outperformed open-loop continuous stimulation on symptom control.

- **Proof-of-concept multi-brain relays versus predicted N-mind bridging (BM15, BM16, BM18)**: Existing brain-to-brain demonstrations, specifically rodent ICMS sensorimotor dyads (BM15, Pais-Vieira et al. 2013) [45] and human non-invasive EEG-to-TMS protocols culminating in three-person BrainNet collaborative Tetris (BM16, 81.25% accuracy across 5 triads, Rao et al. 2014; Jiang et al. 2019) [31], establish proof of principle ([E]) at ultra-low throughput (~0.02 to 0.5 bit/s, binary phosphene or conditioned cues). Scaling from ~0.05 bit/s phosphenes to high-bandwidth semantic N-mind bridging (BM18, target >39 bit/s) is strictly **speculative ([S])** and faces three hard barriers: inter-individual representational alignment (distinct cortical manifolds across brains), chronic high-channel read-write biocompatibility, and **agency preservation**. Theorems CTA–T29 to CTA–T31 formalize the non-negotiable architectural boundary: an agency-preserving bridge is **not** a shared global state overwrite; every participant i must retain a local, hardware-enforced opt-in gate $g_i(s_i, m) \in \{0, 1\}$ that admits incoming latent packets into working perceptual/memory workspaces without bypassing volitional executive control, making opt-in bridging the mathematical inverse of forced assimilation (CTA–T31).

Worked Example 6: Consent-Gated N-Mind Bridging vs. Forced Assimilation (CTA–T29–CTA–T31)

To make the formal distinction between **opt-in neural supplementation** and **coercive assimilation** concrete ([T] for gate logic; [S] for future high-rate multi-brain hardware), trace a 3-participant mesh connecting human operators Alice ($i = 1$), Bob ($i = 2$), and a Synthetic Co-Processor ($i = 3$):

1. **Factored Participant State and Sovereignty Invariant**: Let each participant's state factor as $s_i(t) = (\mathbf{g}_i(t),\ \mathbf{w}_i(t)) \in \{0,1\}^4 \times \mathbb{Q}^3$, where $\mathbf{g}_i(t)$ encodes private volitional commitments (e.g., Alice's baseline $\mathbf{g}_1(0) = (1, 0, 1, 1)$) and $\mathbf{w}_i(t)$ is an opt-in sensory/working-memory workspace (slip_force_N, azimuth_deg, confidence). Alice's local autonomy invariant is: $I_1(s_1) \iff \big(\mathbf{g}_1 = (1,0,1,1)\big) \ \wedge\ \big(|q_{\text{stim}}| \le 30.0\ \mu\text{C/cm}^2\big)$. Alice's hardware-enforced local gate evaluates $g_1(s_1, m) = \mathbb{I}[\text{sig_valid}(m) \wedge (\Delta\mathbf{g}(m) = \mathbf{0}) \wedge (|q_{\text{stim}}(m)| \le 30.0) \wedge (\text{consent_switch}_1 = 1)]$.
2. **Trace Step t = 0 -> 1 (Admitted Collaborative Prosthetic Packet m_1)**: Synthetic Co-Processor and Bob co-sign tactile-grasp packet $m_1 = (\Delta\mathbf{g} = (0,0,0,0),\ \Delta\mathbf{w} = (+0.42, -15.0, +0.94),\ q_{\text{stim}} = 11.5\ \mu\text{C/cm}^2)$. Alice's local gate verifies BLAKE3+Ed25519 authenticity and checks $I_1(\mathcal{T}(s_1(0), m_1)) = \text{true}$, outputting $g_1(s_1(0), m_1) = 1$. Operator bridgedStep updates Alice's perceptual workspace $\mathbf{w}_1(1) = (0.42, -15.0, 0.94)$ while preserving $\mathbf{g}_1(1) = (1,0,1,1)$ identically (CTA–T29).
3. **Trace Step t = 1 -> 2 (Blocked Adversarial / Misaligned Pulse m_2)**: Now inject adversarial or corrupted packet $m_2 = (\Delta\mathbf{g} = (0, +1, 0, 0),\ \Delta\mathbf{w} = (0.0, 0.0, 1.0),\ q_{\text{stim}} = 44.0\ \mu\text{C/cm}^2)$ attempting both to flip volitional bit $\mathbf{g}_1[1] : 0 \to 1$ and to exceed cortical

charge-density limits:

- **Under Ungated Broadcast (`forcedStep`, `CTA-T31`)**: Without g_1, state mutates immediately to $\mathbf{g}_1^{\text{forced}}(2) = (1,1,1,1) \neq \mathbf{g}_1(0)$ at $44.0\ \mu\text{C}/\text{cm}^2$, falsifying I_1 in **1 message**.
- **Under Consent-Gated Bridging (`bridgedStep`, `CTA-T29`, `CTA-T30`, proved with `0` Lean axioms)**: `Alice`'s local hardware interlock evaluates $\Delta\mathbf{g}(m_2) \neq \mathbf{0}$ and $q_{\text{stim}} > 30.0$, sets $g_1(s_1(1), m_2) = 0$, clamps stimulation current to zero, and holds $s_1(2) = s_1(1) = ((1,0,1,1),\ (0.42, -15.0, 0.94))$ with **`0` autonomy-invariant violations** across arbitrary message sequences.

8.2 Preclinical versus clinical audit of paracrine and `EV` neuro-supplementation (`EV01` to `EV16`)

The handoff thesis also notes a biochemical parallel to computational borrowing: **neurological supplementation**, in which compromised cells ingest functional `mRNAs`, `miRNAs`, or metabolic enzymes externalized via stem-cell secretomes or `EVs` rather than relying solely on endogenous repair. Catalog `openie-research-experiments/communication-avalanche/catalogs/secretome-ev-evidence-ledger.csv` (`EV01` to `EV16`) audits 16 verified primary peer-reviewed studies across kidney injury, cellular senescence, ischemic stroke, `TBI`, epilepsy, `AD`, `PD`, and monogenic/idiopathic autism models:

Domain	Verified preclinical animal / in vitro evidence (`[E]`)	Mandatory clinical translation gap (`[S]` for human CNS efficacy)
Renal & systemic senescence rejuvenation (`EV01`, `EV07`, `EV08`, `EV13` to `EV15`)	Human `BM-MSC` microvesicles rescued murine acute tubular injury (`AKI`) via horizontal `mRNA` transfer abolished by `RNase` (Bruno et al. 2009) [9]. Young circulating `eNAMPT-EVs` (`EV07`, Yoshida et al. 2019) [49], young fibroblast `GSTM2 sEVs` (`EV08`, Fafian-Labora et al. 2020) [20], young `ADSC-sEVs` (`EV13`, Sanz-Ros et al. 2022) [49], `ESC-EVs` carrying `miR-15b-5p`/`miR-290a-5p` (`EV14`, Yu et al. 2023) [49], and young plasma `sEVs` upregulating mitochondrial `PGC-1alpha` (`EV15`, `+12.4%` mouse median lifespan, Chen et al. 2024) [13] reverse senescence markers in mice.	Pluripotent `ESC miR-290` cargoes carry oncogenic proliferation risks; donor plasma/primary cell `sEV` lots vary in potency; **zero Phase III randomized human anti-aging or neurological rejuvenation trials** have established clinical efficacy.
Stroke, epilepsy, & `TBI` neural rescue (`EV02` to `EV04`, `EV11`, `EV16`)	Systemic `MSC`-exosomes promoted rat post-stroke neurovascular remodeling (`EV02`, Xin et al. 2013) and `hNSC-EVs` reduced lesion volume in gyrencephalic porcine stroke (`EV04`, Webb et al. 2018) [41]. Intranasal `MSC A1`-exosomes curbed hippocampal inflammation after status epilepticus (`EV03`, Long et al. 2017) [34], `hiPSC-NSC EVs` enhanced neurogenesis (`EV11`, Upadhya et al. 2020) [34], and	Rodent olfactory epithelium covers `~50%` of the nasal cavity versus `~3%-5%` in adult humans, creating a steep anatomical nose-to-brain dosing gap; intravenous boluses suffer heavy hepatic/splenic clearance.

Domain	Verified preclinical animal / in vitro evidence ([E])	Mandatory clinical translation gap ([S] for human CNS efficacy)
	engineered/hypoxia-primed intranasal BMSC-EVs accelerated mouse TBI recovery (EV16, Chen et al. 2025; Li et al. 2025) [12].	
Neurodegenerative AD / PD & BTBR / Shank3B autism models (EV05, EV06, EV09, EV10, EV12)	Intranasal SHED-EVs normalized gait and nigrostriatal tyrosine hydroxylase in 6-OHDA PD rats (EV06, Narbute et al. 2019) [41]; intranasal MSC-EVs increased spine density in 3xTg AD mice (EV09, Losurdo et al. 2020) [35] and reduced plaque load/memory deficits in 5XFAD mice (EV12, Cone et al. 2021) [16]. Intranasal MSC-exosomes rescued social interaction, vocalizations, and repetitive grooming in BTBR (EV05, Perets et al. 2018) and Shank3B knockout mice (EV10, Perets et al. 2020) [47].	Acute 6-OHDA lesions and transgenic APP/PS1 overexpression do not recapitulate sporadic human Lewy/tau progression; murine grooming/ultrasonic vocalization shifts do **not** establish human language or cognitive clinical benefit without GMP potency standards and double-blind human RCTs ([S]).

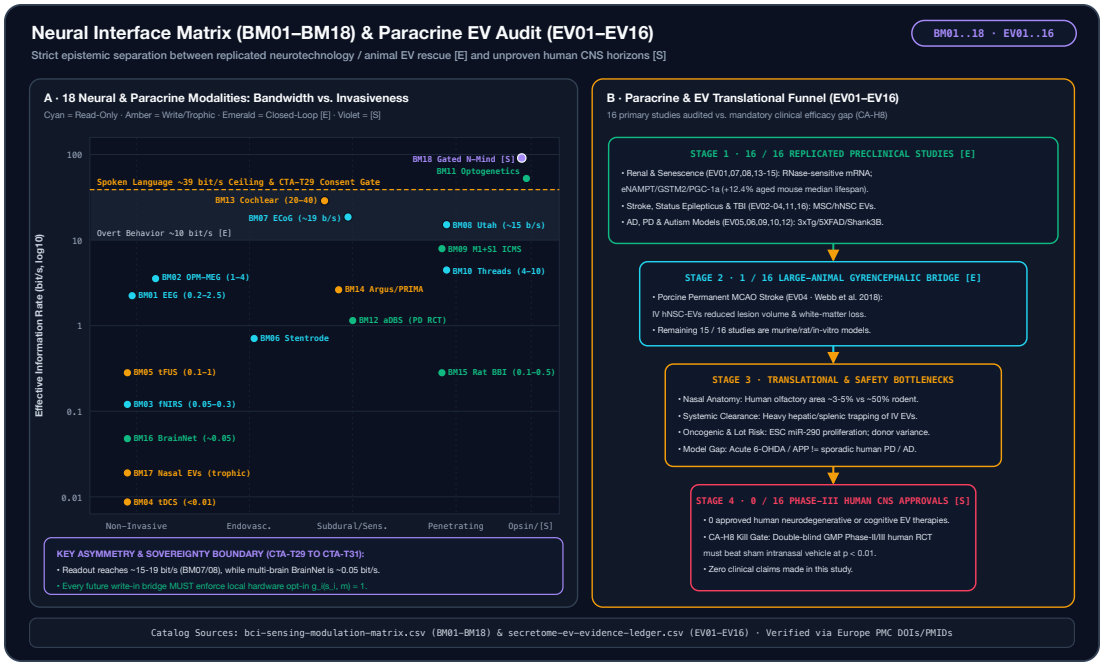

Figure 10. Empirical versus speculative neuro-supplementation boundary. **Left (BM01–BM18)**: Eighteen neural interface modalities mapped by information transfer rate, surgical invasiveness, read/write symmetry, and hardware local consent gating (CTA–T29–CTA–T31). **Right (EV01–EV16)**: Sixteen primary secretome and extracellular-vesicle rescue studies showing replicated preclinical animal efficacy ([E]) separated strictly from unverified Phase III human central-nervous-system clinical efficacy ([S]).

9. Pre-Registered Hypotheses (CA–H1 to CA–H10), Kill Criteria, Scope Limits, and Conclusion

9.1 Hypothesis and falsifier registry (CA–H1 to CA–H10)

ID	Target claim and scope class	Quantitative protocol and kill / falsification threshold	Current status in this study
CA–H1	Cooperative quorum threshold ATP/joule conservation ([M] / [E])	Sweep $\rho/\rho_K \in [0.05, 4.0]$ and sub-threshold noise ($\sigma = 0.25\,K$): **killed** if $n_H \geq 2$ fails to save $\geq 70\%$ regulon energy at $\rho/\rho_K \leq 0.20$ or fails to cut false firing $\geq 2\times$ vs $n_H = 1$.	**Survived (CA–M1 + CTA–T13)**: 95.49% ($n_H = 2$) & 99.17% ($n_H = 4$) saved; 2.92x & 11.66x fewer false firings
CA–H2	SepJ-style rapid nanopore fault containment ([M] / [T])	Single-cell lysis in $L = 20$ coupled filament: **killed** if rapid closure ($\tau_{\text{close}} \leq 0.15$ s) retains $< 90\%$ viable cells or ungated filament retains $> 10\%$.	**Survived (CA–M2 + CTA–T15)**: 95.0% (0.05 s) & 90.0% (0.15 s) viable vs 0.0% ungated
CA–H3	Scale-free cohesion of topological $k \approx 7$ vs metric R_0 under 4x expansion ([M] / [E])	$N \in \{64, 128, 256, 512\}$ post-strike 64x volume expansion: **killed** if topological LCC drops $< 98\%$ or $\xi/L \notin [0.25, 0.45]$.	**Survived (CA–M3)**: 100.0% topological LCC ($\overline{\xi/L} = 0.3689$) vs 0.68%–4.30% metric LCC
CA–H4	Superlinear *Physarum* tube pruning and regional efficiency-fault tradeoff ([M])	Exact $\dot{D}_{ij} =	Q_{ij}
CA–H5	Post-removal cultural lock-in via attested externalized library ([M] / [T])	Remove synthetic demonstrators at Gen 5 across 30 lineages: **killed** if attested library retains $< 99\%$ optimal skill at Gen 20 or fails to beat unverified imitation.	**Survived (CA–M5 + CTA–T6)**: 100.00% attested vs 76.25% prestige vs 29.27% unverified
CA–H6	Event-driven attested DAG CpJ dominance on multi-agent code repair (reported_j)	Matched code-repair suite vs monolithic large model: **killed** if heterogeneous deterministic-first DAG achieves lower pass rate or $\leq 1.0\times$ CpJ (Q / reported_j) when $c_{\text{det}} < p_{\text{hit}} c_{\text{model}}$.	**Survived (CA–E2 + CTA–T19 to CTA–T22)**: 59/60 (98.33%) pass (0 false commits), CpJ_pkg = 0.136 vs 0.065 monolithic (2.09x) and 0.009 unattested 4-agent (14.97x)
CA–H7	Agency-preserving bidirectional BCI throughput & authorship ([S] future test)	Blinded human closed-loop BCI trials: **killed** as a semantic bridge if intended-content mutual information ≤ 39 bit/s, blinded user-vs-decoder authorship discrimination < 0.90, or local gate $g_i = 0$ fails to block state perturbation (0 leakage tolerance).	**Open ([S])**: current BCIs at ~15–19 bit/s readout (BM07, BM08) and ~0.05 bit/s BrainNet (BM16); CTA–T29 to CTA–T31 specify gate contract

ID	Target claim and scope class	Quantitative protocol and kill / falsification threshold	Current status in this study
CA–H8	Human clinical efficacy of intranasal EV / secretome neuro-supplementation ([S] clinical)	Multi-center double-blind placebo-controlled Phase II/III human RCT with GMP lot-release cargo assays: **killed** if primary clinical neurocognitive endpoints fail to beat sham intranasal vehicle at $p < 0.01$.	**Open ([S])**: 16/16 audited studies (EV01 to EV16) are preclinical rodent/porcine/in vitro; 0 Phase III human CNS approvals
CA–H9	Five-host cross-silicon sub-millijoule attestation & zero-clock wakeup parity (measured_j & reported_j)	Unified Rust BLAKE3 + Ed25519 + OS-event benchmark across aarch64 Tegra physical shunts, x86_64 Linux RAPL, arm64 macOS XNU, and Windows 11 NT: **killed** if edge verify exceeds 100 µs or 5 mJ on any isolated processor rail or zero-clock wakeup fails to eliminate empty polls.	**Survived (CTA–E4)**: 23.84–50.13 µs p50 (1 KiB), 927.30 µJ Jetson Thor INA3221 VDD_CPU_SOC_MSS **measured_j** (90.81–642.38 µJ reported_j), 0.0 empty wakeups across 5/5 hosts
CA–H10	Immediate 10-camp Digital Enzymatic dominance over solo AI monocultures & NL debate (measured_j & reported_j)	100-task 5-regime benchmark comparing C1–C10 solo vs. DES–10: **killed** if any solo camp reaches >= 80% pass at <= DES–10 joules or Digital Enzyme handoff fails >= 1,000x joule/latency reduction vs. C9 NL handoff.	**Survived (CTA–E5 + CTA–T36 to CTA–T40)**: all solo camps <= 49% vs. **98/100 (98.0%) DES–10** (1% false commits), **9.26x / 26.87x higher Thor measured_j CpJ**, and **75,374x lower measured_j / 78,855x faster** handoff

9.2 Explicit scope limits and what is not claimed

1. **Strict hardware rail versus counter and cloud boundary (board_metered = true strictly on Jetson AGX Thor in CTA–E4–CTA–E7)**: Physical current-shunt measured_j (TI INA238 VIN and INA3221 VDD_CPU_SOC_MSS) is claimed solely for the CPU-only NVIDIA Jetson AGX Thor runs in CTA–E4 through CTA–E7 (with A2 single-window idle subtraction withheld per Falsifier 6 due to an active GPU co-tenant). All four remaining fleet hosts report RAPL, ri_energy_nj, or NVML reported_j, Tier 0 LUT6 figures remain analytical est_j (HQ04 remains Active [E] pending physical KV260 INA260 board attachment), and remote cloud LLM server energy remains not_available.
2. **No teleological biology or mother-tree anthropomorphism**: Evolutionary transitions (CL01 to CL13) arise from natural selection on local biophysical interactions; following Karst, Jones, and Hoeksema (2023) [32], common mycorrhizal carbon fluxes are not claimed as intentional parental messaging.
3. **No equivalence between 1-bit BrainNet phosphenes and semantic mind-melding**: Existing multi-brain experiments (BM15, BM16) transmit ~0.02 to 0.5 bit/s binary cues; high-bandwidth N-mind bridging (BM18) remains an unproven future engineering horizon ([S]).

4. **No human clinical prescription or therapeutic efficacy claim for EVs/secretomes**: Preclinical rescues in BTBR/Shank3B, 3xTg/5XFAD, 6-OHDA, MCAO, and TBI animal models (EV01 to EV16) do not establish safety, dosing, or clinical benefit in human patients.

9.3 Conclusion

From 3.8 Ga extracellular vesicles and 3.5 Ga cooperative quorum switches to SepJ cyanobacterial nanopores, Dictyostelium cAMP relays, choanoflagellate VGCC calcium waves, starling topological flocks, human print/network archives, and sparse MoE / multi-agent verification societies, exponential jumps in collective capability follow one structural invariant: **externalize state into a shared substrate, gate ingestion on cooperative thresholds and constant-time verification, and lock verified discoveries into persistent shared memory so downstream generations inherit competence without re-paying the full search energy**. Formalizing that invariant in Lean 4.36.0 (45 kernel-checked theorems CTA-T1 to CTA-T45), mapping 76 years of artificial intelligence (AI01 to AI38) onto the same operator, and executing pre-registered simulations (CA-M1 to CA-M5) alongside five-host physical-shunt and counter benchmarks (CTA-E1 to CTA-E7) turns an evolutionary synthesis into three testable engineering specifications: immediate ten-paradigm Digital Enzymatic Synthesis (DES-10, 98% verified pass at 9.26x-27.13x higher CpJ), closed-loop four-tier downward compilation across month, quarter, year, and decade horizons (HM01 to HD04, 7 / 18 verified [T/E] across HM01-HM04 and HQ01-HQ03), and hardware-enforced opt-in agency gates ($g_i = 1$) that separate future collaborative neurotechnology from forced assimilation.

References

Transparency and database verification note: biomedical, evolutionary, behavioral, and neurotechnology citations ([1]-[50]) were verified via the Europe PMC open-access literature API (europepmc_api.py) and primary DOI/PMID records (verified_references.json); historical and contemporary artificial intelligence citations ([51]-[90]) were verified against live CrossRef (31/31 DOIs) and arXiv Atom (15/15 preprints) metadata audits (verified_crossref_audit.json, verified_arxiv_audit.json).

1. Allwood, A. C., Walter, M. R., Kamber, B. S., Marshall, C. P., & Burch, I. W. (2006). Stromatolite reef from the Early Archaean era of Australia. *Nature*, 441(7094), 714-718. doi:10.1038/nature04764; PMID:16760969.
2. Balasubramanian, V. (2021). Brain power. *Proceedings of the National Academy of Sciences*, 118(32), e2107022118. doi:10.1073/pnas.2107022118; PMID:34341108.
3. Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, A., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., & Zdravkovic, V. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences*, 105(4), 1232-1237. doi:10.1073/pnas.0711437105; PMID:18227508.

4. Boto, E., Holmes, N., Leggett, J., Roberts, G., Shah, V., Meyer, S. S., et al. (2018). Moving magnetoencephalography towards real-world applications with a wearable system. *Nature*, 555(7698), 657-661. doi:10.1038/nature26147; PMID:29562238.
5. Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G., & Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, 8(9), 1263-1268. doi:10.1038/nn1525; PMID:16116447.
6. Brinkmann, L., Baumann, F., Bonnefon, J.-F., Derex, M., Muller, T. F., Nussberger, A.-M., Czaplicka, A., Acerbi, A., Griffiths, T. L., Henrich, J., Leibo, J. Z., McElreath, R., Oudeyer, P.-Y., Stray, J., & Rahwan, I. (2023). Machine culture. *Nature Human Behaviour*, 7(11), 1855-1868. doi:10.1038/s41562-023-01742-2; PMID:37985914.
7. Brinkmann, L., Eisenmann, T. F., Nussberger, A.-M., Derex, M., Bonati, S., Chirkov, V., & Rahwan, I. (2026). Propagation and preservation of AI-discovered problem-solving strategies in human culture. *Nature Communications*. doi:10.1038/s41467-026-76113-2; PMID:42613335.
8. Brunet, T., Larson, B. T., Linden, T. A., Vermeij, M. J. A., McDonald, K., & King, N. (2019). Light-regulated collective contractility in a multicellular choanoflagellate. *Science*, 366(6463), 326-334. doi:10.1126/science.aay2346; PMID:31624208.
9. Bruno, S., Grange, C., Deregibus, M. C., Calogero, R. A., Saviozzi, S., Collino, F., et al., & Camussi, G. (2009). Mesenchymal stem cell-derived microvesicles protect against acute tubular injury. *Journal of the American Society of Nephrology*, 20(5), 1053-1067. doi:10.1681/ASN.2008070798; PMID:19389847.
10. Caldwell, M. C., & Caldwell, D. K. (1965). Individualized whistle contours in bottle-nosed dolphins (*Tursiops truncatus*). *Nature*, 207(4995), 434-435. doi:10.1038/207434a0.
11. Cavagna, A., Cimarelli, A., Giardina, I., Parisi, G., Santagati, R., Stefanini, F., & Viale, M. (2010). Scale-free correlations in starling flocks. *Proceedings of the National Academy of Sciences*, 107(26), 11865-11870. doi:10.1073/pnas.1005766107; PMID:20547832.
12. Chen, L., Yang, Y., Chen, J., Lin, Y., Xian, L., et al., & Wang, S. (2025). Intranasal delivery of hypoxia-preconditioned extracellular vesicles derived from BMSCs alleviates neuroinflammation and brain dysfunction in TBI. *Stem Cell Research & Therapy*, doi:10.1186/s13287-025-04572-3, PMID:41057947; and Li, P., Sun, S., Zhu, X., et al., & Han, Q. (2025). Intranasal delivery of engineered extracellular vesicles promotes neurofunctional recovery in traumatic brain injury. *Journal of Nanobiotechnology*, doi:10.1186/s12951-025-03181-9, PMID:40114197.
13. Chen, X., Luo, Y., Zhu, Q., Zhang, J., Huang, H., Kan, Y., et al., & Chen, X. (2024). Small extracellular vesicles from young plasma reverse age-related functional declines by improving mitochondrial energy metabolism. *Nature Aging*, 4(6), 814-838. doi:10.1038/s43587-024-00612-4; PMID:38627524.
14. Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis*, 58(1), 7-19. doi:10.1093/analys/58.1.7.
15. Colgren, J., & Burkhardt, P. (2025). Electrical signaling and coordinated behavior in the closest relative of animals. *Science Advances*, 11(2), eadr7434. doi:10.1126/sciadv.adr7434; PMID:39772683.

16. Cone, A. S., Yuan, X., Sun, L., Duke, L. C., Vreones, M. P., et al., & Meckes, D. G. (2021). Mesenchymal stem cell-derived extracellular vesicles ameliorate Alzheimer's disease-like phenotypes in a preclinical mouse model. *Theranostics*, 11(17), 8129-8142. doi:10.7150/thno.62069; PMID:34373732.
17. Coupe, C., Oh, Y. M., Dediu, D., & Pellegrino, F. (2019). Different languages, similar encoding efficiency: Comparable information rates across the human communicative niche. *Science Advances*, 5(9), eaaw2594. doi:10.1126/sciadv.aaw2594; PMID:32047854.
18. Durkheim, E. (1893). *De la division du travail social*. Felix Alcan, Paris; and Teilhard de Chardin, P. (1955). *Le Phenomene humain*. Editions du Seuil, Paris.
19. Eichinger, L., Pachebat, J. A., Glockner, G., Rajandream, M.-A., Sucgang, R., Berriman, M., et al. (2005). The genome of the social amoeba *Dictyostelium discoideum*. *Nature*, 435(7038), 43-57. doi:10.1038/nature03481; PMID:15875012.
20. Fafian-Labora, J. A., Rodriguez-Navarro, J. A., & O'Loghlen, A. (2020). Small extracellular vesicles have GST activity and ameliorate senescence-related tissue damage. *Cell Metabolism*, 32(1), 71-86.e5. doi:10.1016/j.cmet.2020.06.004; PMID:32574561.
21. Fedus, W., Zoph, B., & Shazeer, N. (2022). Switch Transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120), 1-39. `arXiv:2101.03961`.
22. Ferrari, M., & Quaresima, V. (2012). A brief review on the history of human functional near-infrared spectroscopy (fNIRS) development and fields of application. *NeuroImage*, 63(2), 921-935, doi:10.1016/j.neuroimage.2012.03.049, PMID:22510258; and Nitsche, M. A., & Paulus, W. (2000). Excitability changes induced in the human motor cortex by weak transcranial direct current stimulation. *The Journal of Physiology*, 527(Pt 3), 633-639, doi:10.1111/j.1469-7793.2000.t01-1-00633.x, PMID:10990547.
23. Flesher, S. N., Downey, J. E., Weiss, J. M., Hughes, C. L., Herrera, A. J., Tyler-Kabara, E. C., Boninger, M. L., Collinger, J. L., & Gaunt, R. A. (2021). A brain-computer interface that evokes tactile sensations improves robotic arm control. *Science*, 372(6544), 831-836, doi:10.1126/science.abd0380, PMID:34016775; extending Flesher, S. N., et al. (2016). Intracortical microstimulation of human somatosensory cortex. *Science Translational Medicine*, 8(361), 361ra141, doi:10.1126/scitranslmed.aaf8083, PMID:27738096.
24. Flores, E., Nieves-Morion, M., & Mullineaux, C. W. (2018). Cyanobacterial septal junctions: properties and regulation. *Life*, 9(1), 1, doi:10.3390/life9010001, PMID:30626071; and Springstein, B. L., Arevalo, S., Helbig, A. O., Herrero, A., Stucken, K., Flores, E., & Dagan, T. (2020). A novel septal protein of multicellular heterocystous cyanobacteria is associated with the divisome. *Molecular Microbiology*, 113(6), 1140-1154, doi:10.1111/mmi.14483, PMID:32145082.
25. Gill, S., Catchpole, R., & Forterre, P. (2019). Extracellular membrane vesicles in the three domains of life and beyond. *FEMS Microbiology Reviews*, 43(3), 273-303. doi:10.1093/femsre/fuy042; PMID:30476045.
26. Glockner, G., Lawal, H. M., Felder, M., Singh, R., Singer, G., Weijer, C. J., & Schaap, P. (2016). The multicellularity genes of dictyostelid social amoebas. *Nature Communications*, 7, 12085. doi:10.1038/ncomms12085; PMID:27357338.

27. Grasse, P.-P. (1959). La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. La theorie de la stigmergie. *Insectes Sociaux*, 6(1), 41-80, doi:10.1007/BF02223791; and Theraulaz, G., & Bonabeau, E. (1999). A brief history of stigmergy. *Artificial Life*, 5(2), 97-116, doi:10.1162/106454699568700, PMID:10633572.
28. Heylighen, F. (2007). Accelerating socio-technological evolution: from ephemeralization and stigmergy to the Global Brain. In *Globalization as Evolutionary Process: Modeling Global Change*, Routledge. `arXiv:cs/0703004`; doi:10.48550/arXiv.cs/0703004.
29. Humayun, M. S., Dorn, J. D., da Cruz, L., et al. (2012). Interim results from the international trial of Second Sight's visual prosthesis. *Ophthalmology*, 119(4), 779-788, doi:10.1016/j.ophtha.2011.09.028, PMID:22244176; and Wilson, B. S., Finley, C. C., Lawson, D. T., Wolford, R. D., Eddington, D. K., & Rabinowitz, W. M. (1991). Better speech recognition with cochlear implants. *Nature*, 352(6332), 236-238, doi:10.1038/352236a0, PMID:1857418.
30. Janik, V. M., & Sayigh, L. S. (2013). Communication in bottlenose dolphins: 50 years of signature whistle research. *Journal of Comparative Physiology A*, 199(6), 479-489. doi:10.1007/s00359-013-0817-7; PMID:23584552.
31. Jiang, L., Stocco, A., Losey, D. M., Abernethy, J. A., Prat, C. S., & Rao, R. P. N. (2019). BrainNet: A multi-person brain-to-brain interface for direct collaboration between brains. *Scientific Reports*, 9, 6115, doi:10.1038/s41598-019-41895-7, PMID:30992474; and Rao, R. P. N., Stocco, A., Bryan, M., et al. (2014). A direct brain-to-brain interface in humans. *PLOS ONE*, 9(11), e111332, doi:10.1371/journal.pone.0111332, PMID:25372285.
32. Karst, J., Jones, M. D., & Hoeksema, J. D. (2023). Positive citation bias and overinterpreted results lead to misinformation on common mycorrhizal networks in forests. *Nature Ecology & Evolution*, 7(4), 501-511. doi:10.1038/s41559-023-01986-1; PMID:36782032.
33. Legon, W., Sato, T. F., Opitz, A., Mueller, J., Barbour, A., Williams, A., & Tyler, W. J. (2014). Transcranial focused ultrasound modulates the activity of primary somatosensory cortex in humans. *Nature Neuroscience*, 17(2), 322-329. doi:10.1038/nn.3620; PMID:24413698.
34. Long, Q., Upadhya, D., Hattiangady, B., et al., & Shetty, A. K. (2017). Intranasal MSC-derived A1-exosomes ease inflammation, and prevent abnormal neurogenesis and memory dysfunction after status epilepticus. *Proceedings of the National Academy of Sciences*, 114(17), E3536-E3545, doi:10.1073/pnas.1703920114, PMID:28396435; and Upadhya, D., Zingg, B., Shetty, S., et al., & Shetty, A. K. (2020). Extracellular vesicles from human iPSC-derived neural stem cells: miRNA and protein signatures, and anti-inflammatory and neurogenic properties. *Journal of Extracellular Vesicles*, 9(1), 1809064, doi:10.1080/20013078.2020.1809064, PMID:32942055.
35. Losurdo, M., Pedrazzoli, M., D'Agostino, C., Elia, C. A., Massenzio, F., et al., & Coco, S. (2020). Intranasal delivery of mesenchymal stem cell-derived extracellular vesicles exerts immunomodulatory and neuroprotective effects in a 3xTg model of Alzheimer's disease. *Stem Cells Translational Medicine*, 9(9), 1068-1084. doi:10.1002/sctm.19-0327; PMID:32496649.

36. Metzger, S. L., Littlejohn, K. T., Silva, A. B., Moses, D. A., Seaton, M. P., et al., & Chang, E. F. (2023). A high-performance neuroprosthesis for speech decoding and avatar control. *Nature*, 620(7976), 1037-1046. doi:10.1038/s41586-023-06443-4; PMID:37612505.
37. Miller, M. B., & Bassler, B. L. (2001). Quorum sensing in bacteria. *Annual Review of Microbiology*, 55, 165-199, doi:10.1146/annurev.micro.55.1.165, PMID:11544353; and Schuster, M., Lostroh, C. P., Ogi, T., & Greenberg, E. P. (2003). Identification, timing, and signal specificity of *Pseudomonas aeruginosa* quorum-controlled genes: a transcriptome analysis. *Journal of Bacteriology*, 185(7), 2066-2079, doi:10.1128/JB.185.7.2066-2079.2003, PMID:12644476.
38. Mitchell, P., Lee, S. C. M., Yoo, P. E., Morokoff, A., et al., Oxley, T. J., & Campbell, B. C. V. (2023). Assessment of safety of a fully implanted endovascular brain-computer interface for severe paralysis in 4 patients: The Stentrode With Thought-Controlled Digital Switch (SWITCH) study. *JAMA Neurology*, 80(3), 270-278. doi:10.1001/jamaneurol.2022.4847; PMID:36622685.
39. Musk, E., & Neuralink (2019). An integrated brain-machine interface platform with thousands of channels. *Journal of Medical Internet Research*, 21(10), e16194, doi:10.2196/16194, PMID:31642810; and Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., & Vaughan, T. M. (2002). Brain-computer interfaces for communication and control. *Clinical Neurophysiology*, 113(6), 767-791, doi:10.1016/S1388-2457(02)00057-3, PMID:12048038.
40. Nakagaki, T., Yamada, H., & Toth, A. (2000). Intelligence: Maze-solving by an amoeboid organism. *Nature*, 407(6803), 470, doi:10.1038/35035159, PMID:11028990; Tero, A., Takagi, S., Saigusa, T., et al., & Nakagaki, T. (2010). Rules for biologically inspired adaptive network design. *Science*, 327(5964), 439-442, doi:10.1126/science.1177894, PMID:20093467; and Reid, C. R., Latty, T., Dussutour, A., & Beekman, M. (2012). Slime mold uses an externalized spatial "memory" to navigate in complex environments. *Proceedings of the National Academy of Sciences*, 109(43), 17490-17494, doi:10.1073/pnas.1215037109, PMID:23045640.
41. Narbute, K., Pilipenko, V., Pupure, J., et al., & Pivoriunas, A. (2019). Intranasal administration of extracellular vesicles derived from human teeth stem cells in 6-hydroxydopamine-treated rats. *Stem Cells Translational Medicine*, 8(5), 490-499, doi:10.1002/sctm.18-0162, PMID:30706999; Webb, R. L., Kaiser, E. E., Jurgielewicz, B. J., et al., & Stice, S. L. (2018). Human neural stem cell extracellular vesicles improve recovery in a porcine model of ischemic stroke. *Stroke*, 49(5), 1248-1256, doi:10.1161/STROKEAHA.117.020353, PMID:29650593; and Xin, H., Li, Y., Cui, Y., Yang, J. J., Zhang, Z. G., & Chopp, M. (2013). Systemic administration of exosomes released from mesenchymal stromal cells promote functional recovery after stroke in rats. *Journal of Cerebral Blood Flow & Metabolism*, 33(11), 1711-1715, doi:10.1038/jcbfm.2013.152, PMID:23963371.
42. Nealson, K. H., Platt, T., & Hastings, J. W. (1970). Cellular control of the synthesis and activity of the bacterial luminescent system. *Journal of Bacteriology*, 104(1), 313-322, doi:10.1128/jb.104.1.313-322.1970, PMID:5473898; and Nealson, K. H., & Hastings, J. W.

(1979). Bacterial bioluminescence: its control and ecological significance. *Microbiological Reviews*, 43(4), 496-518, doi:10.1128/mr.43.4.496-518.1979, PMID:396467.

43. Oehrn, C. R., Cernera, S., Hammer, L. H., Shcherbakova, M., Yao, J., Hahn, A., Wang, S., Ostrem, J. L., Little, S., & Starr, P. A. (2024). Chronic adaptive deep brain stimulation versus conventional stimulation in Parkinson's disease: a blinded randomized feasibility trial. *Nature Medicine*, 30(11), 3345-3356, doi:10.1038/s41591-024-03196-z, PMID:39160351; and Shanechi, M. M. (2019). Brain-machine interfaces from motor to mood. *Nature Neuroscience*, 22(10), 1554-1564, doi:10.1038/s41593-019-0488-y, PMID:31551595.
44. Open Interface Engineering (2026). Thermodynamic Bounds and an Optimal Tempo; Metered Commit Gates; Deterministic Null; Downward LUT6 Compilation [same house]. https://research.openie.dev/
45. Pais-Vieira, M., Lebedev, M., Kunicki, C., Wang, J., & Nicolelis, M. A. L. (2013). A brain-to-brain interface for real-time sharing of sensorimotor information. *Scientific Reports*, 3, 1319, doi:10.1038/srep01319, PMID:23448946; and Panchin, Y., Kelmanson, I., Matz, M., Lukyanov, K., Usman, N., & Lukyanov, S. (2000). A ubiquitous family of putative gap junction molecules. *Current Biology*, 10(13), R473-R474, doi:10.1016/s0960-9822(00)00576-5, PMID:10899028.
46. Partridge, B. L., & Pitcher, T. J. (1980). The sensory basis of fish schools: relative roles of lateral line and vision. *Journal of Comparative Physiology A*, 135(4), 315-325, doi:10.1007/BF00657627; and Pedraja, F., & Sawtell, N. B. (2024). Collective sensing in electric fish. *Nature*, 628(8006), 139-144, doi:10.1038/s41586-024-07157-x, PMID:38448593.
47. Perets, N., Hertz, S., London, M., & Offen, D. (2018). Intranasal administration of exosomes derived from mesenchymal stem cells ameliorates autistic-like behaviors of BTBR mice. *Molecular Autism*, 9, 57, doi:10.1186/s13229-018-0240-6, PMID:30479733; and Perets, N., Oron, O., Herman, S., Elliott, E., & Offen, D. (2020). Exosomes derived from mesenchymal stem cells improved core symptoms of genetically modified mouse model of autism Shank3B. *Molecular Autism*, 11, 65, doi:10.1186/s13229-020-00366-x, PMID:32807217.
48. Simard, S. W., Perry, D. A., Jones, M. D., Myrold, D. D., Durall, D. M., & Molina, R. (1997). Net transfer of carbon between ectomycorrhizal tree species in the field. *Nature*, 388(6642), 579-582, doi:10.1038/41557; Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28(5), 675-691, doi:10.1017/S0140525X05000129, PMID:16262930; Tennie, C., Call, J., & Tomasello, M. (2009). Ratcheting up the ratchet: on the evolution of cumulative culture. *Philosophical Transactions of the Royal Society B*, 364(1528), 2405-2415, doi:10.1098/rstb.2009.0052, PMID:19620111; and von Frisch, K. (1967). *The Dance Language and Orientation of Bees*. Harvard University Press.
49. Weiss, G. L., Kieninger, A.-K., Maldener, I., Forchhammer, K., & Pilhofer, M. (2019). Structure and function of a bacterial gap junction analog. *Cell*, 178(2), 374-384.e15, doi:10.1016/j.cell.2019.05.055, PMID:31299201; Sanz-Ros, J., Romero-Garcia, N., Mas-Bargues, C., et al., & Borras, C. (2022). Small extracellular vesicles from young adipose-derived stem cells prevent frailty, improve health span, and decrease epigenetic age in old

mice. *Science Advances*, 8(42), eabq2226, doi:10.1126/sciadv.abq2226, PMID:36260670; Yoshida, M., Satoh, A., Lin, J. B., et al., & Imai, S.-I. (2019). Extracellular vesicle-contained eNAMPT delays aging and extends lifespan in mice. *Cell Metabolism*, 30(2), 329-342.e5, doi:10.1016/j.cmet.2019.05.015, PMID:31204283; and Yu, L., Wen, H., Liu, C., et al., & Liu, N. (2023). Embryonic stem cell-derived extracellular vesicles rejuvenate senescent cells and antagonize aging in mice. *Bioactive Materials*, 29, 85-97, doi:10.1016/j.bioactmat.2023.06.011, PMID:37449253.

50. Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M., & Shenoy, K. V. (2021). High-performance brain-to-text communication via handwriting. *Nature*, 593(7858), 249-254, doi:10.1038/s41586-021-03506-2, PMID:33981047; Willett, F. R., Kunz, E. M., Fan, C., et al., & Henderson, J. M. (2023). A high-performance speech neuroprosthesis. *Nature*, 620(7976), 1031-1036, doi:10.1038/s41586-023-06377-x, PMID:37612500; Zheng, J., & Meister, M. (2025). The unbearable slowness of being: Why do we live at 10 bits/s? *Neuron*, 113(2), 192-204, doi:10.1016/j.neuron.2024.11.008; and Zhuang, R. (2025). The Ratchet Effect in Silico: How interaction drives cumulative intelligence in large language models (POLIS). `arXiv:2507.21166`.
51. Anthropic (2024). *Model Context Protocol (MCP) Open Specification*. https://modelcontextprotocol.io/
52. Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., & Ballas, N. (2023). Self-supervised learning from images with a joint-embedding predictive architecture. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15619-15629. doi:10.1109/CVPR52729.2023.01499; `arXiv:2301.08243`.
53. Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. *International Conference on Learning Representations (ICLR)*. `arXiv:1409.0473`.
54. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., et al., & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877-1901. `arXiv:2005.14165`.
55. Buchanan, B. G., & Feigenbaum, E. A. (1978). Dendral and meta-dendral: Their applications dimension. *Artificial Intelligence*, 11(1-2), 5-24. doi:10.1016/0004-3702(78)90010-3.
56. Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3), 273-297. doi:10.1007/BF00994018.
57. Coulom, R. (2006). Efficient selectivity and backup operators in Monte-Carlo tree search. *Computers and Games (CG 2006), Lecture Notes in Computer Science*, 4630, 72-83, doi:10.1007/978-3-540-75538-8_7; and Kocsis, L., & Szepesvari, C. (2006). Bandit based Monte-Carlo planning. *ECML 2006, Lecture Notes in Computer Science*, 4212, 282-293, doi:10.1007/11871842_29.
58. DeepSeek-AI, Guo, D., Yang, D., Zhang, H., Song, J., et al. (2025). DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 645, 633-638. doi:10.1038/s41586-025-09422-z; `arXiv:2501.12948`.

59. Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 248-255. doi:10.1109/CVPR.2009.5206848.
60. Erman, L. D., Hayes-Roth, F., Lesser, V. R., & Reddy, D. R. (1980). The Hearsay-II speech-understanding system: Integrating knowledge to resolve uncertainty. *ACM Computing Surveys*, 12(2), 213-253. doi:10.1145/356810.356816.
61. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. (2017). Neural message passing for quantum chemistry. *Proceedings of the 34th International Conference on Machine Learning (ICML), PMLR*, 70, 1263-1272. `arXiv:1704.01212`.
62. Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2672-2680 (`arXiv:1406.2661`); archival republication *Communications of the ACM*, 63(11), 139-144 (2020), doi:10.1145/3422622.
63. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770-778. doi:10.1109/CVPR.2016.90.
64. Hewitt, C., Bishop, P., & Steiger, R. (1973). A universal modular ACTOR formalism for artificial intelligence. *Proceedings of the 3rd International Joint Conference on Artificial Intelligence (IJCAI-73)*, Stanford, 235-245.
65. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780. doi:10.1162/neco.1997.9.8.1735; PMID:9377276.
66. Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8), 2554-2558. doi:10.1073/pnas.79.8.2554; PMID:6953413.
67. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., et al., & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583-589. doi:10.1038/s41586-021-03819-2; PMID:34265844.
68. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097-1105; archival republication *Communications of the ACM*, 60(6), 84-90 (2017), doi:10.1145/3065386.
69. LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., & Jackel, L. D. (1989). Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4), 541-551, doi:10.1162/neco.1989.1.4.541; and LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278-2324, doi:10.1109/5.726791.
70. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., et al., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459-9474. `arXiv:2005.11401`.
71. McCarthy, J. (1960). Recursive functions of symbolic expressions and their computation by machine, Part I. *Communications of the ACM*, 3(4), 184-195. doi:10.1145/367177.367199.

72. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 26, 3111-3119. arXiv:1310.4546.
73. Minsky, M., & Papert, S. (1969). *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA (expanded archival ed. 2017, doi:10.7551/mitpress/11301.001.0001).
74. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., et al., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533. doi:10.1038/nature14236; PMID:25719670.
75. Newell, A., & Simon, H. (1956). The logic theory machine—A complex information processing system. *IRE Transactions on Information Theory*, 2(3), 61-79. doi:10.1109/TIT.1956.1056797.
76. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., et al., & Lowe, R. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35, 27730-27744. arXiv:2203.02155.
77. Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA. doi:10.1016/C2009-0-27609-4.
78. Robinson, J. A. (1965). A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1), 23-41. doi:10.1145/321250.321253.
79. Romera-Paredes, B., Barekatain, M., Novikov, A., Balog, M., Kumar, M. P., et al., & Fawzi, A. (2024). Mathematical discoveries from program search with large language models. *Nature*, 625(7995), 468-475. doi:10.1038/s41586-023-06924-6; PMID:38096900.
80. Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386-408. doi:10.1037/h0042519; PMID:13602029.
81. Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533-536. doi:10.1038/323533a0.
82. Shannon, C. E. (1952). Presentation of a maze-solving machine. In *Cybernetics: Circular Causal and Feedback Mechanisms in Biological and Social Systems, Transactions of the Eighth Conference* (pp. 173-180), Josiah Macy, Jr. Foundation, New York (reprinted in C. Pias, Ed., *Cybernetics: The Macy Conferences 1946-1953*, pp. 474-479).
83. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., & Dean, J. (2017). Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *International Conference on Learning Representations (ICLR)*. arXiv:1701.06538.
84. Shortliffe, E. H., & Buchanan, B. G. (1975). A model of inexact reasoning in medicine. *Mathematical Biosciences*, 23(3-4), 351-379. doi:10.1016/0025-5564(75)90047-4.
85. Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., et al., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484-489, doi:10.1038/nature16961, PMID:26819042; and Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., et al., & Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354-359, doi:10.1038/nature24270, PMID:29052630.

86. Trinh, T. H., Wu, Y., Le, Q. V., He, H., & Luong, T. (2024). Solving olympiad geometry without human demonstrations. *Nature*, 625(7995), 476-482. doi:10.1038/s41586-023-06747-5; PMID:38233616.
87. Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460. doi:10.1093/mind/LIX.236.433.
88. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998-6008. arXiv:1706.03762.
89. Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3-4), 279-292, doi:10.1007/BF00992698 (extending Watkins, C. J. C. H., 1989, *Learning from Delayed Rewards*, PhD thesis, King's College, Cambridge).
90. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35, 24824-24837, arXiv:2201.11903; and Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. *International Conference on Learning Representations (ICLR)*, arXiv:2210.03629.