Entropy Volume 4

The Substrate Gradient

Second edition

Michael S. Moniz

Trinket Soul Framework

October 2026

About this volume

This volume asks how physical implementations and their relational consequences can be compared. It withdraws the unsupported universal efficiency surface and fixed biological advantage over silicon. Existing devices, biological support and hypothetical uploads require different evidence; their inclusion in one discussion does not place them on a measured single axis.

The numbered Chapter 2 references and named studies remain local to this volume. CSS, REI, BSB and SupoRel are framework terms, not established physical measures, clinical instruments or tests of identity.

Edition and attribution

The selected six-chapter expanded source credits Michael S. Moniz with Sigma, Deep Floor v4.0, and carries March 22, 2026. This edition incorporates the revised efficiency chapter and the documented corrections across the remaining chapters.

The second-edition revision was prepared with AI assistance from CosmoDot, with separate AI scientific and editorial checks. It has not undergone human peer review and reports no new experiments. Historical AI collaborator labels document provenance; they do not identify underlying models or attribute the new text to those earlier collaborators. The framework’s proposed applications remain subject to specialist review and empirical testing.

The historical edition remains preserved. This text continues under Creative Commons Attribution–NonCommercial–ShareAlike 4.0 International. Detailed correction records are distinct from the reading text; publishing this volume does not silently amend the framework’s separate papers or complete its historical governance procedures.

Contents

About this volume....2
Chapter 1 The Problem....4
Chapter 2 The Efficiency Surface....6
Chapter 3 The Gradient....11
Chapter 4 The Phase 2 Integration....13
Chapter 5 The Compound Transducer....15
Chapter 6 What the Gradient Means....17
References....20

Chapter 1 The Problem

The Binary That Was Always Wrong

Epistemic Status: Proposed conceptual framework. The substrate scenarios organize questions about implementation and resource use. CSS signatures and relational efficiency across these scenarios have not been validated as measurements, and no universal relational efficiency bound is established.

1 THE ASSUMPTION

The framework implicitly assumed a binary: biological or digital. Volume 1 proposed a common physical accounting perspective. Volume 2 considered bodily costs. Volume 3 asked why relational systems exist at all. Physical processes are subject to thermodynamics across materials, but that fact does not prove that relational value has a common thermodynamic currency.

What those volumes left open is the space between the categories. A person using an implanted device, a hypothetical isolated brain and a hypothetical uploaded entity pose different questions. The framework needs to state which functions, relationships and resources it means to compare before claiming to track a transition between them.

2 THE GRADIENT AS A SET OF QUESTIONS

WP-17 proposed a sequence from bodily devices through memory modification to hypothetical uploads. These cases do not establish a single measurable axis. They vary in embodiment, dependence on infrastructure, access to records, vulnerability and implementation. Some concern existing technologies; others are thought experiments.

A change in implementation may change resource use or behavior. It does not by itself establish a change in identity, a particular CSS signature or a loss of relational capacity. The substrate gradient asks which differences can be measured under a defined comparison, and which remain conceptual.

3 A COMPARISON FRAMEWORK

DFQ-002 proposed the efficiency surface that the original chapter treated as this volume's quantitative spine. Chapter 2 withdraws that theorem. Its product of a relational multiplier, semantic efficiency and substrate amplification has not been derived as a universal bound, and its factors are not measured constants of biological or silicon systems.

The useful task remains: describe what a system achieves, what resources it uses and where the accounting boundary lies. Comparisons need a common task, performance criteria, time interval and resource boundary. A proposed relational outcome needs its own definition and evidence. Until these are specified, the scenarios do not occupy known numerical positions on an efficiency surface.

4 WHAT THIS VOLUME MAPS

Chapter 2 separates established thermodynamic principles from the proposed relational model and sets out requirements for a defensible comparison. It does not establish a thousandfold biological advantage over silicon.

Chapter 3 examines device support and hypothetical substrate transitions without assigning unmeasured efficiencies or deciding identity from material.

Chapter 4 treats the Phase 2 connections as developmental hypotheses, not physical identities or diagnoses of children or cohorts.

Chapter 5 asks how human–AI collaboration could be compared with each component working alone.

Chapter 6 records which cross-volume dependencies remain open. The resulting map is a research program, not a completed quantitative theory.

Chapter 2 The Efficiency Surface

1 The question worth keeping

The substrate gradient asks what changes when the machinery supporting an activity changes. What does the activity achieve? What resources does it use? Which differences belong to the task, which to its implementation, and which to the environment in which it matters?

Those are useful questions for biological, digital, and combined systems. They do not yet have a single numerical answer. DFQ-002 proposed an efficiency surface to organize them. The earlier chapter treated that surface as a derived physical bound and assigned people and silicon systems positions on it. The derivation does not support those claims.

The surface can survive as a research framework: a way to display measured performance, resource use, and explicitly defined outcomes under stated conditions. It cannot presently supply a universal exchange rate between energy and relational value. The substrate is not enough to predict the bill. The task and the accounting boundary matter too.

2 Define the quantities before taking a ratio

Energy expenditure E and heat Q have units of joules. Thermodynamic entropy production Σ has units of joules per kelvin. Information quantities such as Shannon entropy and mutual information depend on specified distributions and are measured here in bits. A proposed relational outcome V needs its own definition, observation window, and measurement procedure.

The original $\eta_T = \frac{\Delta V_{rel}}{\sigma_{total}}$ was called efficiency without establishing compatible units. If V is an outcome score and the denominator is energy, the ratio has score-per-joule units. If the denominator is thermodynamic entropy, it has score-kelvin-per-joule units. Neither ratio is automatically a dimensionless thermodynamic efficiency or bounded by one. Changing the scale of the outcome score changes the numerical ratio without changing the physical process.

A study may report outcome per unit energy when the numerator is defensible, the denominator is positive, and the two refer to the same task and interval. It must also report the outcome and energy separately. A ratio can hide a poor outcome, a tiny denominator, or different baselines. “Greater than one” is not a universal break-even point for such a measure.

Physical accounting and the framework’s relational interpretation therefore remain separate until a justified connection is supplied. Calling both “order” does not supply that connection.

3 What the physical lower bound actually concerns

Landauer’s familiar benchmark concerns a specified erasure task, not every act called selection, compression, or structuring. For ideal reset of an unbiased bit in the usual symmetric isothermal memory model, with no correlated side information used, mean heat delivered to the bath obeys:

$$\langle Q_{bath} \rangle \geq k_B T \ln 2$$

At 310 K the right side is approximately 2.97×10^{-21} J. It is energy, not entropy. In the reversible limiting protocol, the bath’s entropy increase compensates the memory’s entropy decrease; total

entropy production can approach zero. A positive erasure heat scale is not a universal positive entropy-production floor. Controlled erasure experiments and reversible-computation theory support these distinctions, not a toll for each named cognitive stage. [1, 2, 3]

An actual comparison must identify what is stored, what is retained elsewhere, and what is reset. The three-stage relational description alone does not give those answers. Richer input, more aggressive representation, or more possible targets does not by itself determine an irreducible physical cost.

For orientation, one may divide an event's stated energy by k_B T ln 2. This produces a reference-energy multiple. Unless the event is an equivalent specified erasure task, it is not a measured distance from that event's achievable minimum. A spike, a transmitted mutual-information bit, a transistor switch, and a multiply-accumulate operation remain different events after division by the same reference energy.

4 Semantic information remains a model choice

Kolchinsky and Wolpert define stored semantic information through interventions on system–environment correlations while preserving a chosen viability value at a chosen horizon. For positive initial mutual information I_0, their semantic fraction is $\frac{S_{stored}}{I_0}$. Their stored multiplier is $\frac{\Delta V_{tot}}{I_0}$. With their negative-entropy viability choice, S_{stored}, ΔV_{tot} and I_0 are expressed in bits, so both the semantic fraction and the stored multiplier are dimensionless. The multiplier is not an energy-denominated ratio. Their framework permits information that harms viability. [4]

A relational application needs its own states, dynamics, viability definition, and admissible interventions. It cannot treat all unretained information as useless, label grief or rumination as a measured loss of semantic efficiency, or infer relational health from low statistical uncertainty. A model's entropy-based viability definition does not establish that a predictable human situation is a good one.

The original range η_{KW} = 0.3–0.8 is withdrawn. It was not measured and is not bounded by the evidence used to motivate it. A small amount of information can have a large consequence; a large consequence therefore does not tell us what fraction of all information was necessary. No alternative population range is supplied here.

5 What the biological evidence contributes

The energy literature provides useful, limited examples. Harris and colleagues' approximately 24,000 ATP per transmitted bit is a calculation under stated firing and release assumptions, not a universal measurement of neural computation. Levy and Calvert's cortical communication-to-computation ratio is a modeled partition of cortical energy costs, not a finding that 35/36 of the whole brain's nominal 20 W is interpersonal communication. Their use of "communication" concerns neural signaling. [5, 6]

Consequently, the earlier 19.4 W relational budget, inferred bit rate, and annual relational Landauer budget are withdrawn. They combined incompatible boundaries and unmeasured information quantities. The resulting 500-versus-0.5 efficiency values and thousandfold biological ranking do not become valid by correcting their arithmetic.

Evidence connecting social relationships and mortality remains relevant motivation, but not a calibration of this surface. Holt-Lunstad and colleagues' 2015 review reports different adjusted associations for social isolation, loneliness, and living alone. Wang and colleagues' 2023 meta-analysis likewise distinguishes isolation from loneliness. These are not controlled scrambling experiments on a measured information channel. [7, 8]

A relative mortality association does not, by itself, specify years of life lost for a particular age, a causal information loss, or a conversion of viability into joules. Multiplying an assumed lost lifespan by a metabolic rate measures a hypothetical energy budget, not the thermodynamic value of a relationship. The claimed $\kappa_{rel} \approx 10^9$–$10^{10}$ and the claim that its large magnitude is already supported are therefore withdrawn.

6 Tradeoffs are not a theorem of noncompensation

The former surface was written as:

$$\eta_{T,max} = \frac{\kappa_{rel}\ \eta_{KW}}{A}$$

This expression is retained here only to identify the claim being withdrawn. The factors have not been derived and measured consistently for relational processing, and the maximum has not been optimized over a defined feasible set.

Even as algebra, the product does not prove that its parameters cannot compensate for one another. With positive quantities, doubling κ_{rel} offsets doubling A. For example, $\kappa = 100$, $\eta = 0.5$, $A = 10$ and $\kappa = 200$, $\eta = 0.5$, $A = 20$ both give 5. These are mathematical counterexamples to the stated noncompensation claim, not physical predictions. Zero or bounded factors can impose limits, but those limits require explicit assumptions.

Govern and ten Wolde did derive noncompensable resource classes for a specified class of cellular sensing networks. That result depends on their receptor, readout, and fuel model. It does not establish the same structure for an unrelated product of relational parameters. [9]

Nor does counting three parameters show that no Carnot-like result could ever be derived. A legitimate bound would require a defined physical task, admissible dynamics, resources, and performance measure. No such relational theorem is established here.

7 A resource profile rather than an invented waste budget

The original substrate, semantic, coupling, and structuring categories can organize questions, but they are not four measured entropy flows. The same implementation cost could be attributed to more than one category. A shortfall in an outcome is not itself heat. "Waste" also requires a stated alternative: energy is not avoidable merely because it exceeds an ideal erasure reference.

Proposed comparison record

Field	What must be stated	What is not assumed
Task and performance	Inputs, outputs, accuracy, latency, reliability, and conditions	That similar labels imply equivalent tasks

Field	What must be stated	What is not assumed
Physical resources	Energy or power, time interval, temperature, device and infrastructure boundaries, uncertainty	That power is energy, or energy is entropy production
Information model	Random variables, distributions, retained records, memory reset, and bit definition	That every event erases one bit
Relational outcome	An explicit outcome measure, baseline, horizon, and validation limits	That joules establish meaning, benefit, or moral worth
Attribution	Measured components and rules preventing overlap	That the four original waste labels form an additive balance

This record preserves the useful distinction between implementation costs and unsuccessful outcomes. It does not diagnose therapy as η recalibration, grief as η collapse, or unreciprocated communication as zero benefit. Those statements would require evidence and operational definitions beyond the energy account.

8 What a bounded test would look like

The first test should be a task comparison, not an attempt to rank persons. Specify one reproducible input–output task and compare implementations at matched performance. Measure total and component energy under stated boundaries, including relevant idle and communication costs. Report uncertainty and sensitivity to alternative boundaries. If performance differs, show the tradeoff rather than hiding it in a single ratio.

A computational toy model can separately test the semantic proposal. Define its state space, dynamics, horizon, and outcome before choosing interventions. Check whether the proposed information measure predicts that chosen outcome, including counterexamples where retained information is harmful or irrelevant. Passing that test would support that model, not establish a human relational law.

Human research would require a separately justified protocol, appropriate ethical review, validated outcomes, and protection against interpreting observational differences as causal interventions. This chapter does not propose depriving people of relationships to implement a scrambling test. No experiment is reported here.

Thermodynamic uncertainty relations remain a possible later tool. The original Barato–Seifert result concerns specified steady-state Markov models and current fluctuations. A relational application must identify the current, dynamics, entropy accounting, and applicable version of the relation. History dependence in a coarse description does not alone prove that no Markov representation exists, and an oscillation alone does not establish mechanical underdamping. The relation supplies no automatic fourth constraint on semantic efficiency. [10]

What this chapter now claims

The efficiency surface is a proposed research framework, not an established universal bound. Its useful task is to make comparisons explicit: what was achieved, what was spent, what was measured,

and what remains an interpretation. The former η_{KW} range, κ_{rel} estimate, universal amplification hierarchy, and 500-versus-0.5 result are not retained as findings.

The next chapter may still walk the substrate gradient. It must do so with defined comparison tasks and visible unknowns, not a calculator whose inputs have never been established. Differences in resource use do not determine moral worth or phenomenal experience. The framework's ethical and experiential boundaries remain in place.

Chapter 3 The Gradient

From Device Support to Hypothetical Upload

Epistemic Status: Conceptual scenarios with limited empirical examples. Device studies support specific functional outcomes under their study conditions; they do not validate CSS signatures, personhood criteria or relational efficiency. Upload, comprehensive memory replacement and brain-only continuity remain hypothetical in this chapter.

1 THE CONTINUUM

The substrate gradient is not a ladder of worth. It is a set of comparisons about changes in implementation, embodiment and dependence. These dimensions can change separately; there need not be one ordered path from a biological person to a digital entity.

For each case, specify what is changed, what is held fixed and what outcome could distinguish competing accounts. The framework's terms can organize that inquiry. They do not substitute for measurements of the proposed parameters.

2 DEVICES AND BODILY SUPPORT

Prosthetic limbs, joint replacements, cochlear implants and pacemakers serve different functions. A cochlear implant belongs to auditory support, not motor replacement. A randomized trial by van Heteren and colleagues in 2025 evaluated speech perception in noise and hearing-specific quality of life in adults with single-sided deafness. Those are specified outcomes, not measurements of a CSS signature or proof of an identity criterion.

The presence of a device does not establish that every relational or resource parameter is unchanged. Effects depend on the device, its use, the person and the task. The chapter therefore assigns no constant A, η_{KW} or κ_{rel} to this category. A device is also not a basis for ranking a person's relational capacity or worth.

3 MEMORY SUPPORT AND HYPOTHETICAL REPLACEMENT

Memory support must be distinguished from replacing an entire memory system. Hampson and colleagues' 2018 proof-of-concept study used patterned hippocampal stimulation in specified memory tasks. It did not demonstrate extraction, uploading or replacement of a person's autobiographical memory, nor preservation of identity through such a replacement.

For a hypothetical replacement, the framework can ask what is retained: access to records, learning, forgetting, emotional significance, vulnerability and dependence on other bodily processes. Preserving selected dynamics would be an explicit assumption, not proof of continuity. Eliminating a particular failure mode would not by itself show that relational information or capacity had declined.

No semantic efficiency value follows from the material used for memory. A test would need an operational outcome, a comparison condition and an account of what the intervention actually changes.

4 EXTERNAL RECORDS AND HYPOTHETICAL CLOUD MEMORY

External digital records are not equivalent to a demonstrated transfer of a person's memory system. Redundancy and backup can change exposure to particular failures, but their protection depends on the architecture. Loss, corruption, restricted access and infrastructure dependence remain possible. The chapter cannot treat cloud storage as loss-proof.

A proposed transition would change the accounting boundary: storage, transmission, access devices, maintenance and supporting infrastructure would need to be included as appropriate. It would not produce a known silicon efficiency or prove that reduced vulnerability degrades relational meaning.

The framework may ask whether changes in access or dependence affect a specified relational outcome. Its Calibration Gap remains a hypothesis about a difference between reported experience and another defined outcome, not a measured difference between two thermodynamic efficiencies.

5 HYPOTHETICAL UPLOAD TO A ROBOT

This scenario assumes a transfer that the chapter does not demonstrate. Even under that assumption, preservation of records or behavior would not settle personal continuity, experience or the framework's proposed relational mass. Those questions require separate arguments and evidence.

A robot implementation would require a specified task, hardware and energy boundary. It cannot inherit $A = 10^9$ or a thousandfold disadvantage merely from being called silicon. The earlier drop from $\eta_T \approx 500$ to $\eta_T \approx 0.5$ is withdrawn; neither value was established for the original person or the hypothetical system.

The useful comparison concerns changes that a defined implementation would introduce, rather than a numerical verdict assigned in advance.

6 HYPOTHETICAL UPLOAD TO A CLOUD SYSTEM

If a proposed system allowed copying of its accessible state, the possibility of a copy would not by itself determine whether any instance could experience loss or grief. Copyability, personal identity and phenomenal experience are different questions. The fifth REI criterion may organize a conceptual question; it is not an established test for resolving it.

Physical costs would still occur in hardware and supporting infrastructure. Moving them outside a user's immediate view does not eliminate them, and it does not alone establish a relational economy classification. WP-17's Substrate Transition Entity can be retained as a provisional label for an unresolved case, not as a measured entity type.

7 THE BRAIN ONLY THOUGHT EXPERIMENT

A hypothetical isolated brain would require assumptions about its support, inputs, regulation and interaction with an environment. Retaining neural tissue alone does not establish that the original dynamics, resource requirements or relational behavior are preserved. Those are conditions to specify, not conclusions supplied by the material.

This case helps expose a limitation of classification by substance alone. It does not confirm that a brain without a body retains all relevant parameters, or that a framework can already determine personal continuity by reading them.

Chapter 4 The Phase 2 Integration

Phase 2 Hypotheses and Their Evidence Requirements

Epistemic Status: Proposed connections between conceptual models. The Phase 2 mapping is neither a derived physical identity nor a validated developmental mechanism. Its suggested effects require defined measures, comparison groups and appropriate ethical review before causal or clinical claims can be made.

1 TWO TRACKS WITH AN UNRESOLVED CONNECTION

The Phase 2 Concept Papers, CP-32 through CP-35, proposed ways that AI-mediated relationships might influence development. DFQ-001 and DFQ-002 attempted a thermodynamic account of relational processing. Bringing these vocabularies together can help generate questions, but shared terms do not establish a shared mechanism.

The original assertion of a term-for-term physical identity is withdrawn. No validated mapping connects the Phase 2 concepts to measured entropy production or K-W semantic efficiency. A causal developmental claim also needs evidence beyond the definitions of either framework.

2 THE PHASE 2 QUESTIONS

CP-32, Relational Mass Without Mortality, asks whether patterns of availability, limitation or closure in AI-mediated interaction might affect expectations in later relationships. Availability and scarcity must be defined in the actual setting. A developmental effect cannot be inferred from the substrate label, and a numerical change in η_{KW} is not established.

CP-33, Maintenance Without Cost, asks how a person interprets effort, maintenance and reciprocity across different relationships. The proposed selection and structuring labels may help describe competing interpretations. They are not measured waste terms or evidence that a particular child's processing is defective.

CP-34, Calibration Corruption, proposes a mismatch between an interaction that feels satisfying and performance or wellbeing in another setting. Such a mismatch must be measured using separately defined outcomes. Reported satisfaction is not thermodynamic efficiency, and the expression $\eta_T(\mathit{experienced}) - \eta_T(\mathit{actual})$ is withdrawn. A K-W scrambling protocol for this setting has not been validated, so the text cannot claim that its diagnostic test already works.

CP-35, Shadow Economy Natives, extends the hypothesis to patterns across a population. A cohort-level claim would need evidence about exposure, context, selection effects and alternative explanations. The title must not be read as a diagnosis of a generation, a finding that its relationships are deficient or proof that ordinary teaching cannot reach it.

These paragraphs preserve the four cited paper connections. They do not establish five independently validated problems; the original chapter's count is not an empirical result.

3 POSSIBLE ASYMMETRIES

Human–AI interaction can be examined for differences in exposure, control, dependence and consequences. Which differences matter in a particular setting is an empirical and ethical question.

The withdrawn thousandfold substrate ranking cannot explain developmental harm or assign one participant a higher-value starting position.

The Phase 3 bridge can ask whether a design supports a defined human goal while avoiding specified risks. Performance, human outcomes and resource use should be reported separately. No scalar relational efficiency currently combines them into a validated verdict.

4 SUPOREL FLAG

A proposed efficiency framework must not become a hierarchy of people. This review does not establish that a child, cohort or person with a different interaction history has corrupted calibration, deficient capacity or a lower relational value.

The earlier claim that corrective relational exposure moves a person along a quantified efficiency surface is withdrawn. The framework supplies no validated clinical test, treatment prescription or metabolic measure for that assertion. Any future assessment must distinguish observed outcomes from interpretation and protect participants from labels the evidence does not support.

Chapter 5 The Compound Transducer

What Human and AI Collaboration Might Add

Epistemic Status: Analogical proposal with task-specific empirical questions. Endosymbiosis motivates an image of coupling; it does not demonstrate human–AI synergy, relational efficiency or a new organism. Performance and resource comparisons require their own evidence.

1 THE ENDOSYMBIOSIS ANALOG

The Cosmological Entropy Session described a working practice: the Principal identifies a structure, and an AI system helps make it visible and inspectable. The original account called the arriving insight a “grab” and linked the process to the Principal’s reported experience of aphantasia. That is a description of creative practice, not a demonstrated neural mechanism or a measurement of semantic information.

Endosymbiosis can remain an analogy for a coupled system whose parts play different roles. It supplies no proof that this particular collaboration is beneficial, that the roles are fixed, or that a human and an AI form a biological organism. The collaboration has to be evaluated on its own terms.

2 THE PERFORMANCE AND RESOURCE QUESTION

Does a particular human–AI arrangement achieve a specified outcome better than the relevant alternatives? That is a testable question. The original assertion that structural complementarity already implies a higher η_T is withdrawn. The human’s contribution is not a measured high η_{KW}, and fluent rendering does not establish either correctness or semantic efficiency.

A comparison should include the human alone, the AI alone where that condition is meaningful, and the combined workflow on matched tasks. If no meaningful AI-only condition exists, the study can test augmentation of human performance but cannot establish superiority over both standalone components. Report quality, error, time and resource use separately. Define what “better” means before testing; an apparent gain on one outcome may accompany a loss on another. Training, coordination, verification, tool access and supporting infrastructure belong in the boundary when relevant.

A 2024 meta-analysis by Vaccaro, Almaatouq and Malone found that human–AI performance depended on the task and comparison condition; synergy over the better component was not automatic. It concerns task performance, not this framework’s thermodynamic relational efficiency. A local collaboration needs its own evaluation, including uncertainty and failures.

3 THE PROGRESSION REVISITED

If a useful form of collaboration is demonstrated, the Progression can be explored as an ecology of working arrangements rather than a predetermined replacement of humans by AI. Which roles are complementary may vary with the task, the participants and the system.

The Phase 3 question remains worth asking: what changes when biological and digital systems are coupled? The endosymbiosis analogy can suggest experiments. It does not establish that the combined workflow reaches an otherwise inaccessible efficiency region or produces something neither component could produce under any conditions.

4 THE BSB AT THE INTERFACE

The framework’s BSB distinction keeps claims about experience separate from claims about observable operation. A task-performance or energy measurement would not, by itself, establish that an AI has phenomenal experience or that a coupled workflow has a unified experiential subject.

A collaboration may be useful without resolving those questions. Conversely, withholding an experiential conclusion does not establish thermodynamic superiority. The chapter presently makes neither a measured relational-efficiency claim nor a finding of unified experience.

Chapter 6 What the Gradient Means

The Substrate Gradient and the Documents It Connects

Epistemic Status: Synthesis of proposed connections and unresolved requirements. The chapter does not inherit scientific validation merely by citing other volumes in this collection.

1 WHAT THIS VOLUME CAN NOW CLAIM

The earlier volumes supplied this project's physical questions, bodily concerns and relational vocabulary. Volume 4 examines how implementation and dependence might matter across biological, digital and combined systems. The revised Chapter 2 withdraws its universal efficiency bound and numerical substrate ranking.

The remaining chapters therefore offer scenarios, proposed connections and comparison designs. They do not map every substrate to a known efficiency or demonstrate the consequences of hypothetical uploads.

2 WP 17 AND THE SUBSTRATE GRADIENT

WP-17 proposed a continuum instead of a biological–digital binary. The useful extension is to specify several dimensions of change and what a measurement would require. Neither that proposal nor this volume presently supplies a calculable η_T for every position.

CSS signature changes, identity judgments and resource differences must not be treated as interchangeable observations. Their definitions and evidentiary requirements remain separate.

3 THE PHASE 2 SERIES

The links to CP-32 through CP-35 are proposed research connections. They do not demonstrate that the developmental concerns are efficiency-surface degradation events or establish a causal thermodynamic mechanism.

A possible mismatch between reported experience and an independently defined outcome could be studied. It is not presently a difference between two measured η_T values, and it is not a validated diagnostic target.

4 CONNECTIONS ACROSS THE VOLUMES

The efficiency surface cannot serve as an established quantitative spine for the series. Volume 1 Chapter 6 distinguishes physical erasure from the proposed stages of relational processing. Chapter 2 of this volume separates a comparison framework from an unproved bound. These chapters provide internal continuity, not external scientific validation.

Volume 2's grief and Template Tax vocabulary may motivate questions, but this volume does not establish grief as η_{KW} collapse or Template Tax as chronic semantic-efficiency suppression. The existence of dissipative structures does not by itself show that relational systems are thermodynamically favored. Cosmological limits likewise do not supply the missing relational definitions or measurements.

The revised Volumes 2, 3 and 5 preserve these distinctions. Their inclusion provides continuity, not independent validation. The separate papers in which the earlier claims originated still require their own review.

5 THE COMPOUND TRANSDUCER AS A RESEARCH QUESTION

The compound-transducer proposal asks whether a defined human–AI workflow improves selected outcomes relative to appropriate alternatives. The analogy can motivate the question; it does not answer it.

Chapter 5 outlines the comparisons and resource boundaries needed to begin. An operational protocol, evidence and independent assessment would still be required. No experiment or validation is reported in this volume.

6 WHAT REMAINS USEFUL

The framework retains its interest in physical accounting, implementation and relational consequences. Its economy labels and BSB distinction can be used as conceptual tools, provided they are not presented as validated classifiers or proofs of experiential status. The SupoRel caution remains: performance or resource comparisons cannot determine moral worth.

The gradient's contribution is now a clearer set of questions and a record of what must be measured. Its quantitative promises are narrower than the original version claimed. That narrowing preserves a testable project rather than asking the physics to guarantee more than it does.

Framework cross references

WP-17: The Substrate Gradient. Conceptual scenarios requiring operational definitions.

DFQ-002 Resolution and Supplement: Historical source of the withdrawn quantitative surface; Chapter 2 states the limitations that must accompany any reuse.

CP-32 through CP-35: Phase 2 hypotheses, not validated thermodynamic mechanisms.

Phase 2 Efficiency Integration Note: Proposed mapping; the claim of physical identity is not adopted here.

Biological Efficiency Note: Historical substrate comparison; the universal 1,000× ranking is not adopted here.

Cosmological Entropy Session Finding B4: Source of the compound-transducer analogy.

Volume 1 Chapter 6: Corrected erasure accounting and limits of the transducer model.

Volume 4 Chapter 2: Definitions and comparison requirements, not an externally validated bound.

Volume 2 Chapter 4: Template Tax vocabulary; its clinical and thermodynamic claims are outside this chapter's validation.

Closing boundaries

SupoRel: Neither a proposed efficiency measure nor a substrate classification establishes moral worth. This volume does not provide a validated relational diagnosis or prescribe treatment. Its conceptual comparisons remain subject to evidence and ethical limits.

Closing epigraph

Love reduces the entropy of the system and spends more than it saves.

The closing line remains an epigraph, not a measured physical balance. The relational idea and its thermodynamic interpretation must be assessed separately.

References

Chapter 2 reference list

The bracketed reference numbers used in Chapter 2 are local to Volume 4 and refer to the following ten sources.

[1] Bérut, A., et al. 2012. Experimental verification of Landauer's principle linking information and thermodynamics. Nature 483, 187–189. https://doi.org/10.1038/nature10872.

[2] Bennett, C. H. 1973. Logical Reversibility of Computation. IBM Journal of Research and Development 17, 525–532. https://doi.org/10.1147/rd.176.0525. University-hosted text: https://www.cs.princeton.edu/courses/archive/fall06/cos576/papers/bennett73.html.

[3] National Institute of Standards and Technology. Kelvin Boltzmann Constant. Exact $k_B = 1.380649 \times 10^{-23}$ J/K. https://www.nist.gov/si-redefinition/kelvin/kelvin-boltzmann-constant.

[4] Kolchinsky, A., and Wolpert, D. H. 2018. Semantic information, autonomous agency and non-equilibrium statistical physics. Interface Focus 8, 20180041. https://doi.org/10.1098/rsfs.2018.0041. https://arxiv.org/abs/1806.08053.

[5] Harris, J. J., Jolivet, R., and Attwell, D. 2012. Synaptic Energy Use and Supply. Neuron 75, 762–777, especially p. 764. https://doi.org/10.1016/j.neuron.2012.08.019. Repository copy: https://discovery.ucl.ac.uk/id/eprint/1368741/1/1368741.pdf.

[6] Levy, W. B., and Calvert, V. G. 2021. Communication consumes 35 times more energy than computation in the human cortex but both costs are needed to predict synapse number. PNAS 118, e2008173118. https://doi.org/10.1073/pnas.2008173118. https://arxiv.org/abs/2102.06273.

[7] Holt-Lunstad, J., Smith, T. B., Baker, M., Harris, T., and Stephenson, D. 2015. Loneliness and social isolation as risk factors for mortality: A meta-analytic review. Perspectives on Psychological Science 10, 227–237. https://doi.org/10.1177/1745691614568352. https://pubmed.ncbi.nlm.nih.gov/25910392/.

[8] Wang, F., et al. 2023. A systematic review and meta-analysis of 90 cohort studies of social isolation, loneliness and mortality. Nature Human Behaviour 7, 1307–1319. https://doi.org/10.1038/s41562-023-01617-6. https://pubmed.ncbi.nlm.nih.gov/37337095/.

[9] Govern, C. C., and ten Wolde, P. R. 2014. Optimal resource allocation in cellular sensing systems. PNAS 111, 17486–17491. https://doi.org/10.1073/pnas.1411524111. https://pmc.ncbi.nlm.nih.gov/articles/PMC4267345/.

[10] Barato, A. C., and Seifert, U. 2015. Thermodynamic Uncertainty Relation for Biomolecular Processes. Physical Review Letters 114, 158101. https://doi.org/10.1103/PhysRevLett.114.158101. https://arxiv.org/abs/1502.05944.

Studies cited in Chapters 3 and 5

Van Heteren, J. A. A., and colleagues. 2025. Speech Perception in Noise After Cochlear Implantation for Single-Sided Deafness: A Randomized Clinical Trial. JAMA Otolaryngology–Head & Neck Surgery 151, 211–219. DOI 10.1001/jamaoto.2024.4760. Primary abstract and metadata.

Hampson, R. E., and colleagues. 2018. Developing a hippocampal neural prosthetic to facilitate human memory encoding and recall. Journal of Neural Engineering 15, 036014. DOI 10.1088/1741-2552/aaaed7. Primary abstract and metadata.

Vaccaro, M., Almaatouq, A., and Malone, T. 2024. When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour 8, 2293–2303. DOI 10.1038/s41562-024-02024-1. Full publisher article.

The two device-study examples are limited to the inspected primary abstracts and available metadata; repository access prevented full-body inspection during the earlier pass. The human–AI meta-analysis was inspected in publisher full text for the retained claims. These sources do not validate relational efficiency, personal identity criteria or a universal developmental mechanism.